

A General Theory of Concave Regularization for High Dimensional Sparse Estimation Problems

Cun-Hui Zhang*
 Department of Statistics
 and Biostatistics
 Rutgers University, NJ
czhang@stat.rutgers.edu

Tong Zhang†
 Department of Statistics
 and Biostatistics
 Rutgers University, NJ
tzhang@stat.rutgers.edu

Abstract

Concave regularization methods provide natural procedures for sparse recovery. However, they are difficult to analyze in the high dimensional setting. Only recently a few sparse recovery results have been established for some specific local solutions obtained via specialized numerical procedures. Still, the fundamental relationship between these solutions such as whether they are identical or their relationship to the global minimizer of the underlying nonconvex formulation is unknown. The current paper fills this conceptual gap by presenting a general theoretical framework showing that under appropriate conditions, the global solution of nonconvex regularization leads to desirable recovery performance; moreover, under suitable conditions, the global solution corresponds to the unique sparse local solution, which can be obtained via different numerical procedures. Under this unified framework, we present an overview of existing results and discuss their connections. The unified view of this work leads to a more satisfactory treatment of concave high dimensional sparse estimation procedures, and serves as guideline for developing further numerical procedures for concave regularization.

1 Introduction

Let \mathbf{X} be an $n \times p$ design matrix and $\mathbf{y} \in \mathbb{R}^n$ a response vector satisfying

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (1)$$

where $\boldsymbol{\beta} \in \mathbb{R}^p$ is a target vector of regression coefficients and $\boldsymbol{\varepsilon} \in \mathbb{R}^n$ is a noise vector. This paper concerns the estimation of the value of $\mathbf{X}\boldsymbol{\beta}$, that of $\boldsymbol{\beta}$, or its support set $\text{supp}(\boldsymbol{\beta})$, where $\text{supp}(\mathbf{b}) := \{j : b_j \neq 0\}$ for any vector $\mathbf{b} = (b_1, \dots, b_p)^\top \in \mathbb{R}^p$.

We are interested in the high-dimensional case where n and p are both allowed to diverge, including the case of $p \gg n$. We assume that the target vector $\boldsymbol{\beta}$ is sparse in some sense; such as the ℓ_0 sparsity $|\text{supp}(\boldsymbol{\beta})| \leq c_0 n / \ln p$, or the capped- ℓ_1 sparsity $\sum_{j=1}^p \min(1, |\beta_j|/\sigma) \leq c_0 n / \ln p$, where σ is a certain noise level and c_0 is a fixed small constant. While we are mainly interested in

*Research partially supported by the NSF Grants DMS 0906420, DMS-11-06753 and NSA Grant H98230-11-1-0205

†Research partially supported by the following grants: AFOSR-10097389, NSA -AMS 081024, NSF DMS-1007527, and NSF IIS-1016061

the Gaussian noise $\varepsilon \sim N(0, \sigma^2 \mathbf{I}_{n \times n})$ or zero-mean sub-Gaussian noise, the specific noise properties required in our analysis will be provided later.

We consider the following class of penalized least squares estimators

$$\hat{\boldsymbol{\beta}} := \arg \min_{\mathbf{b} \in \mathbb{R}^p} L_\lambda(\mathbf{b}), \quad L_\lambda(\mathbf{b}) := \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\mathbf{b}\|_2^2 + \sum_{j=1}^p \rho(b_j; \lambda), \quad (2)$$

where $\mathbf{b} = (b_1, \dots, b_p)^\top$ and $\rho(t; \lambda)$ is a scalar regularization function with a certain regularization parameter $\lambda > 0$. As an example, we may let $\rho(t; \lambda) = \lambda^2 I(t \neq 0)/2$, which corresponds to the ℓ_0 regularization problem. Here $I(\cdot)$ denotes $\{0, 1\}$ valued indicator function. Since $I(t \neq 0)$ is a discontinuous function at $t = 0$, the corresponding ℓ_0 optimization problem may be difficult to solve. In practice, one also looks at continuous regularizers that approximate ℓ_0 regularization, such as $\rho(t; \lambda) = \min(\lambda^2/2, \lambda|t|)$. As we will show in the paper, sparse local solutions of such regularizers can be obtained using standard numerical procedures (such as gradient descent), and they are closely related to the global solution of (2).

2 Survey of Existing Concave Regularization Results

While this survey is not intended to be comprehensive, it presents a high-level view of some important contributions to the area of concave regularization. We will discuss both methodological and analytical contributions.

2.1 Terminologies

The following notation is used throughout the paper. For any dimension d , bold face letters denote vectors and normal face their elements, e.g. $\mathbf{v} = (v_1, \dots, v_d)^\top$, with $\text{supp}(\mathbf{v})$ being its support $\{j : v_j \neq 0\} \cap \{0, \dots, d\}$. Capital bold face letters denote matrices, e.g. \mathbf{X} and $\boldsymbol{\Sigma}$. The ℓ_q “norm” of \mathbf{v} is $\|\mathbf{v}\|_q := (\sum_{j=1}^d |v_j|^q)^{1/q}$ for $0 < q < \infty$, with the usual extension $\|\mathbf{v}\|_0 := |\text{supp}(\mathbf{v})|$ and $\|\mathbf{v}\|_\infty := \max_{j \leq d} |v_j|$. Design vectors, or columns of \mathbf{X} , are denoted by \mathbf{x}_j . For simplicity, we assume throughout the paper that the columns \mathbf{X} are normalized to

$$\|\mathbf{x}_j\|_2 = \sqrt{n}.$$

This condition is not essential but it simplifies some notations. For variable sets $A \subseteq \{1, \dots, p\}$, $\mathbf{X}_A = (\mathbf{x}_j, j \in A)$ denotes the restriction of columns of \mathbf{X} to A , and $\mathbf{b}_A = (b_j, j \in A)^\top$ the restriction of vector $\mathbf{b} \in \mathbb{R}^p$ to A . The maximum and minimum eigenvalues of matrix $\boldsymbol{\Sigma}$ are denoted by $\lambda_{\max}(\boldsymbol{\Sigma})$ and $\lambda_{\min}(\boldsymbol{\Sigma})$.

Definition 1. *The following terminologies will be used to simplify discussion.*

- (a) *The ℓ_0 sparsity of $\boldsymbol{\beta}$ means $\|\boldsymbol{\beta}\|_0 \leq s^*$. To allow $\boldsymbol{\beta}$ with many more components near zero, a weaker notion of capped- ℓ_1 sparsity is $\sum_j \min(1, |\beta_j|/\lambda_{\text{univ}}) \leq s^*$, where $\lambda_{\text{univ}} = \sigma \sqrt{(2/n) \ln p}$ is the universal threshold level for a certain noise level σ .*
- (b) *A regularity condition on \mathbf{X} is a class \mathcal{X} of (column-normalized) matrices that match a sparsity condition on $\boldsymbol{\beta}$ to guarantee a desired result. Such a regularity condition can be stated as $\mathbf{X} \in \mathcal{X}_{s^*}^{n \times p}$, with matrix classes $\mathcal{X}_{s^*}^{n \times p} \subseteq \mathbb{R}^{n \times p}$ indexed by (n, p, s^*) , where s^* is the*

sparsity level of the matching regularity condition on β . Such a condition on \mathbf{X} is called an ℓ_2 regularity condition (or simply ℓ_2 regular) if the matrix classes $\mathcal{X}_{s^*}^{n \times p}$ are sufficiently large to satisfy the following condition:

- Given any $u_0 \geq 1$, there exists a constant $c_0 > 0$ such that for all $0 < \delta \leq 1/e$

$$\inf_{\mu, n, p, s^*} \left\{ \mu(Q^{-1}(\mathcal{X}_{s^*}^{n \times p})) : \mu \in \mathcal{M}_{u_0}^{n \times p}, (s^*/n) \ln(p/\delta) \leq c_0, \min(n, p, s^*) \geq 1 \right\} \geq 1 - \delta,$$

where $\mathcal{M}_{u_0}^{n \times p}$ is the set of probability measures in $\mathbb{R}^{n \times p}$ under which the rows of $\mathbb{R}^{n \times p}$ are iid $N(0, \Sigma)$ for some Σ with $\lambda_{\max}(\Sigma)/\lambda_{\min}(\Sigma) \leq u_0$ and identical diagonal elements, and Q is the column normalization mapping given by $Q(\mathbf{X}) = (\mathbf{x}_j n^{1/2} / \|\mathbf{x}_j\|_2, j \leq p)$.

- (c) An estimator $\hat{\beta}$ is selection consistent if $\text{supp}(\hat{\beta}) = \text{supp}(\beta)$, and sign-consistent if $\text{sgn}(\hat{\beta}) = \text{sgn}(\beta)$, with the convention $\text{sgn}(0) = 0$ for the sign function.
- (d) An estimator has the oracle property if

$$\hat{\beta} = \hat{\beta}^o, \quad \hat{\beta}_S^o = (\mathbf{X}_S^\top \mathbf{X}_S)^{-1} \mathbf{X}_S^\top \mathbf{y}, \quad \text{supp}(\hat{\beta}^o) \subseteq S, \quad (3)$$

where $S = \text{supp}(\beta)$. The estimator $\hat{\beta}^o$ is called the oracle LSE.

Remark 1. The standard regularity condition for the classical low-dimensional statistical scenario of $p \leq n$ is that the rank of \mathbf{X} is p . Definition 1 (b) generalizes this classical regularity condition to allow $p \gg n$. We may explicitly include the classical situation into the definition of ℓ_2 regularity (that is, require $\mathcal{X}_{s^*}^{n \times p}$ to contain all column-normalized $n \times p$ matrices of rank p) if we confine our discussion to fixed sample conditions. See the last paragraph of this subsection for more discussion.

Remark 2. If we consider a sequence of models in (1) with $n \rightarrow \infty$, then asymptotically an estimator has the oracle property (allowing statistical inference for all linear functionals of β) if

$$\sup_{\mathbf{a}} P\{|\mathbf{a}^\top (\hat{\beta} - \hat{\beta}^o)|^2 > \epsilon \text{Var}(\mathbf{a}^\top \hat{\beta}^o)\} = o(1) \quad \forall \epsilon > 0,$$

and this is a weaker requirement than (3) because it allows $\hat{\beta}$ to converge only asymptotically to $\hat{\beta}^o$. While this work focuses on the stronger requirement (3) that is easier to interpret in the finite sample situation, the weaker definition has been used in some previous asymptotic analysis.

For $0 < r \leq 1$, the capped- ℓ_1 sparsity condition holds for all vectors with $\|\beta\|_r \leq R$ as long as $(R/\lambda_{\text{univ}})^r \leq s^*$.

In the classical statistical scenario of $p < n$, a standard regularity condition on the design matrix \mathbf{X} is that the rank of \mathbf{X} is p . Definition 1 (b) generalizes this classical regularity condition to $p \gg n$. For example, $\inf_{|A| \leq 3s^*} \{\text{rank}(\mathbf{X}_A)/|A|\} = 1$ is ℓ_2 regular. The ℓ_2 notion allows an assessment of the strength of assumptions on \mathbf{X} by random matrix theory without repeating technical statements of more specialized conditions. Moreover, since the ℓ_2 criterion is required to hold for $(s^*/n) \ln(p/\delta) \leq c_0$, results based on ℓ_2 regularity condition on \mathbf{X} and matching sparsity condition of β must apply to the case of large p , including $p \gg n$. Since regularity conditions on β and \mathbf{X} must work together to guarantee their consequences, for simplicity the sparsity level s^* for ℓ_2 regularity is always understood in the sequel as the ℓ_0 or capped- ℓ_1 sparsity level of β given in Definition 1 (a).

Throughout the paper, \mathbf{X} and $\boldsymbol{\beta}$ in (1) are treated as deterministic. Since the ℓ_2 criterion is about the size of $\mathcal{X}_{s^*}^{n \times p}$, it does not imply randomness of \mathbf{X} . In fact, since the ℓ_2 criterion is required to hold simultaneously for all $\mu \in \mathcal{M}_{u_0}^{n \times p}$ with the same $\mathcal{X}_{s^*}^{n \times p}$ in $\mathbb{R}^{n \times p}$, an ℓ_2 regularity condition is weaker than the condition of a random \mathbf{X} with distribution $\mu(Q^{-1}(\cdot))$ for a fixed $\mu \in \mathcal{M}_{u_0}^{n \times p}$ and typically requires a more explicit specification of the matrix class $\mathcal{X}_{s^*}^{n \times p}$. We call the criterion ℓ_2 , since it depends only on the range of the spectrum (the smallest and largest eigenvalues) of $\boldsymbol{\Sigma}$.

The rest of the subsection discusses different forms of ℓ_2 conditions. Since the meaning of sparsity level is always clear in its proper context, for simplicity we will discuss design matrix conditions without explicitly referring to their sparsity levels.

In what follows, we will briefly explain some ℓ_2 -regularity conditions appeared in the literature. Related conditions have been introduced first in the compressive sensing literature to analyze ℓ_1 -regularized recovery of a sparse $\boldsymbol{\beta}$ from its random projection $\mathbf{X}\boldsymbol{\beta}$ with iid $N(0, 1)$ entries in \mathbf{X} . The most well-known of such conditions is the restricted isometry condition (RIP) introduced in [10]. In order to explain RIP, we first define the lower and upper sparse eigenvalues as

$$\kappa_-(m) := \min_{\|\mathbf{u}\|_0 \leq m; \|\mathbf{u}\|_2 = 1} \|\mathbf{X}\mathbf{u}\|_2^2/n, \quad \kappa_+(m) := \max_{\|\mathbf{u}\|_0 \leq m; \|\mathbf{u}\|_2 = 1} \|\mathbf{X}\mathbf{u}\|_2^2/n. \quad (4)$$

RIP requires $\delta_k + \delta_{2k} + \delta_{3k} < 1$ with $k = \|\boldsymbol{\beta}\|_0$ and $\delta_m = \max\{\kappa_+(m) - 1, 1 - \kappa_-(m)\}$. A related condition is the uniform uncertainty principle (UUP) $\delta_{2k} + \theta_{2k,k} < 1$ in [9], where $\theta_{k,\ell} = \max(\mathbf{X}_A \mathbf{v}_A)^\top (\mathbf{X}_B \mathbf{u}_B)/n$ with $A \cap B = \emptyset$, $|A| = k$, $|B| = \ell$, and $\|\mathbf{u}\|_2 = \|\mathbf{v}\|_2 = 1$. For ℓ_1 regularized estimators, bounds of the optimal order for the ℓ_2 -norm estimation error $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2$ can be obtained under RIP, UUP, as well as their improvement $\delta_{1.25k} + \theta_{1.25k,k} < 1$ in [7]. While the conditions for RIP and UUP are specialized to hold for random designs with covariance matrix $\boldsymbol{\Sigma} = \mathbf{I}_{p \times p}$, related conditions using sparse eigenvalues can be defined to fulfill the ℓ_2 criterion in Definition 1(b); for example the sparse Riesz condition (SRC) $\|\boldsymbol{\beta}\|_0 < \max_m 2m/\{1 + \kappa_+(m)/\kappa_-(m)\}$ in [43, 42], and some other extensions in [44, 41]. These more general conditions are ℓ_2 regularity conditions by our definition, and they lead to ℓ_2 -norm estimation error bounds of the optimal order for ℓ_1 regularized estimators. Additional refinements were introduced in the literature, such as the restricted eigenvalue of [3, 22],

$$\text{RE}_2 = \text{RE}_2(\xi, S) := \inf_{\mathbf{u}} \left\{ \|\mathbf{X}\mathbf{u}\|_2 / (\|\mathbf{u}\|_2 n^{1/2}) : \|\mathbf{u}_{S^c}\|_1 < \xi \|\mathbf{u}_S\|_1 \right\}$$

where $S = \text{supp}(\boldsymbol{\beta})$, and the compatibility factor of [37, 39],

$$\text{RE}_1 = \text{RE}_1(\xi, S) := \inf_{\mathbf{u}} \left\{ |S|^{1/2} \|\mathbf{X}\mathbf{u}\|_2 / (\|\mathbf{u}_S\|_1 n^{1/2}) : \|\mathbf{u}_{S^c}\|_1 < \xi \|\mathbf{u}_S\|_1 \right\}.$$

It can be shown that $\text{RE}_1 \geq \text{RE}_2$ and appropriate sparse eigenvalues imply $\text{RE}_2 > 0$. Therefore both RE_2 and RE_1 are ℓ_2 regularity conditions. Moreover, for ℓ_1 regularized estimators, RE_1 provides ℓ_1 -norm estimation and ℓ_2 -norm prediction error bounds of optimal order, and RE_2 provides ℓ_2 -norm estimation bounds of optimal order.

This paper employs an even weaker condition involving a restricted invertibility factor RIF_q in (14) which is related to the cone invertibility factor CIF_q ($q \geq 1$) defined below:

$$\text{CIF}_q = \text{CIF}_q(\xi, S) := \inf \left\{ \frac{|S|^{1/q} \|\mathbf{X}^\top \mathbf{X}\mathbf{u}\|_\infty}{n \|\mathbf{u}\|_q} : \|\mathbf{u}_{S^c}\|_1 < \xi \|\mathbf{u}_S\|_1 \right\}. \quad (5)$$

The quantity CIF_q and its sign-restricted version have appeared in [41], where invertibility factor-based ℓ_q error bounds of the form (19) below have been proven to sharpen earlier results for the Lasso and Dantzig selector [9, 43, 3, 44, 39] when $q \in [1, 2]$. Such error bounds are of optimal order [41, 30]. Of special interests are $q \in [1, 2]$ for which the condition $\text{CIF}_q > 0$ on \mathbf{X} is ℓ_2 regular and

$$\text{CIF}_1(\xi, S) \geq \frac{\text{RE}_1^2(\xi, S)}{(1 + \xi)^2}, \quad \text{CIF}_2(\xi, S) \geq \frac{\text{RE}_1(\xi, S)\text{RE}_2(\xi, S)}{(1 + \xi)} \geq \frac{\text{RE}_2^2(\xi, S)}{(1 + \xi)}. \quad (6)$$

Thus, $\text{CIF}_q > 0$ is an ℓ_2 regularity condition for $q \in [1, 2]$.

A main advantage of using invertibility factor is that for $q > 2$, invertibility factors still yield ℓ_q error bounds of optimal order which match results in [44, 41]. However, the sparse and restricted eigenvalues do not yield error bounds of optimal order due to the unboundedness of $\max_{\|\mathbf{u}\|_2=1} \|\mathbf{u}_S\|_q \|\mathbf{u}_S\|_1 / |S|^{1/q}$ in $|S|$.

We shall point out that different ℓ_2 regularity conditions are typically not equivalent since different norms are involved in the definitions of different quantities. For instance, in a specific example given in [3, 39], RE_1 and CIF_2 , uniformly bounded from away from zero, yield ℓ_1 and ℓ_2 error bounds of optimal order respectively, but RE_2 does not.

In the above discussion, we focus on fixed sample conditions like $\text{RE}_2 > 0$ and $\text{CIF}_2 > 0$, which hold when $\text{rank}(\mathbf{X}) = p$. These conditions can be directly seen as ℓ_2 regular from their existing lower bounds for $p > n$ such as those in [3, 41]. The optimality of the order of the error bounds based on such quantities can be also stated as ℓ_2 regularity conditions by comparing them with sparse eigenvalues. See Remark 7 for more discussion.

2.2 Previous Results

Among concave penalties, the ℓ_1 penalty is the only convex one. Thus, the Lasso (ℓ_1 regularization) is a special case of (2) with $\rho(t; \lambda) = \lambda|t|$ [35, 11]:

$$\hat{\boldsymbol{\beta}}^{(\ell_1)} = \arg \min_{\mathbf{b} \in \mathbb{R}^p} \left[\frac{1}{2n} \|\mathbf{X}\mathbf{b} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{b}\|_1 \right]. \quad (7)$$

As a function of λ , the Lasso path $\hat{\boldsymbol{\beta}} = \hat{\boldsymbol{\beta}}^{(\ell_1)}(\lambda)$ matches that of ℓ_1 constrained quadratic programming. One may use the homotopy/Lars algorithm to compute the complete Lasso path for $\lambda \in [0, \infty)$ [28, 29, 13] or simply use a standard convex optimization algorithm to compute the Lasso solution for a finite set of λ . The Dantzig selector, proposed in [9], is an ℓ_1 -minimization method related to the Lasso, which solves

$$\hat{\boldsymbol{\beta}} = \arg \min_{\mathbf{b} \in \mathbb{R}^p} \|\mathbf{b}\|_1 \quad \text{subject to } \|\mathbf{X}^\top(\mathbf{X}\mathbf{b} - \mathbf{y})\|_\infty \leq \lambda.$$

It has analytical properties similar to that of Lasso, but can be computed by linear programming rather than quadratic programming as in Lasso. Analytic properties of the Lasso or Dantzig selector have been studied in [21, 17, 26, 36, 48, 40, 9, 6, 38, 43, 27, 3, 22, 44, 39, 7, 41]. A basic story is described in the following two paragraphs.

Under various ℓ_2 regularity conditions on \mathbf{X} and the ℓ_0 sparsity condition on $\boldsymbol{\beta}$, the Lasso and Dantzig selector control the estimation errors and the dimension of the selected model in the sense

$$\frac{\|\mathbf{X}\hat{\boldsymbol{\beta}} - \mathbf{X}\boldsymbol{\beta}\|_2^2}{\sigma^2 \ln p} + \frac{\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_q^q}{\{(\sigma^2/n) \ln p\}^{q/2}} + \|\hat{\boldsymbol{\beta}}\|_0 = O_P(s^*), \quad 1 \leq q \leq 2, \quad (8)$$

[9, 38, 43, 3, 22, 44, 41, 5]. Compared with the oracle $\widehat{\beta}^o$ in (3), the estimation loss of $\widehat{\beta}$ is inflated by a factor of no greater order than $\sqrt{\ln p}$, and the size of the selected model is of the same order as the true one. When $\ln(p/n) \asymp \ln p$, it has been proved in [41, 30] that (8) matches the order of the risk of a Bayes estimator for a class of (weak) signals close to zero, so that the order of this loss inflation factor $\sqrt{\ln p}$ is the smallest possible without further assumption on the strength of the signal β . This inflation factor can be viewed as the cost of not knowing $\text{supp}(\beta)$. Nevertheless, when β is strong (in the sense that its minimum nonzero coefficient is not close to zero), then it is possible to achieve the oracle property, which removes the inflation factor. However even in such cases, the logarithmic inflation is still present for the Lasso solution, and it is generally referred to as the *Lasso bias*; it means that the Lasso does not have the oracle property even when the signal is strong [14, 15]. Nonconvex penalty can be used to remedy this issue. For the Lasso and Dantzig selector, extensions of (8) have been established for capped- ℓ_1 sparse β [43, 44, 41] and for $2 < q \leq \infty$ under certain ℓ_q regularity conditions on \mathbf{X} [44, 41]. Error bounds of type (8) have been used in the analysis of the joint estimation of the noise level $\sigma^* := \|\varepsilon\|_2/\sqrt{n}$ and β [32, 2, 33, 34]. For example, the scaled Lasso

$$\{\widehat{\beta}, \widehat{\sigma}\} = \arg \min_{\{\beta, \sigma\}} \left\{ \|\mathbf{y} - \mathbf{X}\beta\|^2 / (2n\sigma) + 2\sqrt{(\ln p)/n} \|\beta\|_1 \right\}$$

provides $|\widehat{\sigma}/\sigma^* - 1| = O_P(|S|(\ln p)/n)$ along with (8) under ℓ_2 regularity conditions [34].

For variable selection, the Lasso is sign consistent in the event

$$\text{sgn}(\widehat{\beta}^o) = \text{sgn}(\beta), \quad \min_{j \in S} |\widehat{\beta}_j^o| \geq \theta_1^* \lambda, \quad \lambda \geq \frac{\sigma \sqrt{(2/n) \ln(p - |S|)}}{(1 - \theta_2^*)_+}, \quad (9)$$

where $\theta_1^* = \|(\mathbf{X}_S^\top \mathbf{X}_S/n)^{-1} \text{sgn}(\beta_S)\|_\infty$, $\theta_2^* = \|\mathbf{X}_{S^c}^\top \mathbf{X}_S (\mathbf{X}_S^\top \mathbf{X}_S)^{-1} \text{sgn}(\beta_S)\|_\infty$, $S = \text{supp}(\beta)$, and $\widehat{\beta}^o$ is the oracle estimator in (3) [26, 36, 48, 40]. Since $\|\widehat{\beta}^o - \beta\|_\infty = O_P(1) \sqrt{(\ln \|\beta\|_0)/n} = o_P(\lambda)$ under mild conditions, θ_1^* and θ_2^* are key quantities in (9). For fixed $\kappa_0 < 1$, $\theta_2^* \leq \kappa_0$ is called the neighborhood stability/strong irrepresentable condition [26, 48]. For \mathbf{X} with iid $N(0, \Sigma)$ rows and given S , θ_1^* and θ_2^* are within a small fraction of their population versions with Σ in place of $\mathbf{X}^\top \mathbf{X}/n$ [40]. For random β with $\|\beta\|_0 \lesssim n/\{\|\mathbf{X} \mathbf{X}^\top/p\|_2 \ln p\}$ and uniformly distributed $\text{sgn}(\beta)$ given $\|\beta\|_0$, $\theta_1^* \leq 2$ and $\theta_2^* \leq 1 - 1/\sqrt{2}$ with large probability under the incoherence condition $\max_{j \neq k} |\mathbf{x}_j^\top \mathbf{x}_k/n| \lesssim 1/(\ln p)$ [8]. It is worth mentioning that neither the incoherence condition nor the strong irrepresentable condition is ℓ_2 regular: in fact they may both fail with $\theta_2^* \asymp |S|^{1/2}$ and $\min_{j \in S} |\widehat{\beta}_j^o| \geq \theta_1^* \lambda$ even in the classical setting of \mathbf{X} being rank p . Since $\theta_2^* \leq 1$ is necessary for the selection consistency of the Lasso under the first two conditions of (9) [36, 40], this means that Lasso is not model selection consistent under ℓ_2 regularity conditions. In order to achieve model selection consistency under ℓ_2 regularity, we have to employ a nonconvex penalty in (2).

For sparse estimation, ℓ_0 penalized LSE corresponds to the choice of $\rho(t; \lambda) = \lambda^2/2I(t \neq 0)$ in (2), and it was introduced in the literature [1, 24, 31] before Lasso. Formally,

$$\widehat{\beta}^{(\ell_0)} = \arg \min_{\beta \in \mathbb{R}^p} \left[\frac{1}{2n} \|\mathbf{X}\beta - \mathbf{y}\|_2^2 + \frac{\lambda^2}{2} \|\beta\|_0 \right]. \quad (10)$$

This method is important for sparse recovery because with the Gaussian noise model $\varepsilon \sim N(0, \sigma^2 I)$, uniform distribution on support set, and flat distribution of β within support, it is a Bayesian procedure for support set recovery. However, this penalty is not easy to work with numerically

because it is discontinuous at zero. The Lasso can be viewed as a convex surrogate of (10), but it does not achieve model selection consistency under ℓ_2 regularity, nor does it have the oracle property when the signal is uniformly strong.

Continuous concave penalties other than Lasso have been introduced to remedy these problems. These concave functions approximate ℓ_0 penalty better than Lasso, and thus can remove the Lasso bias problem. Most concave penalties are interpolations between the Lasso and the ℓ_0 penalty. For example the ℓ_α (bridge) penalty [16] with $0 < \alpha < 1$ is equivalent to the choice of $\rho(t; \lambda) = |t|^\alpha \lambda^{2-\alpha} \{2(1-\alpha)\}^{1-\alpha} / (2-\alpha)^{2-\alpha}$ in (2). While the bridge penalty is continuous, its derivative is ∞ at $t = 0$, which may still cause numerical problems. In fact, the ∞ derivative value means that $\hat{\beta} = 0$ is always a local solution of (2) for bridge penalty, which prevents any possibility for the uniqueness of a reasonable local solution among sparse local solutions—a topic which we will investigate in this paper. In order to address this issue, additional penalty functions $\rho(t; \lambda)$ with finite derivatives at $t = 0$ have been suggested in the literature, such as the SCAD penalty [14], and the MCP penalty [42]. These penalties can be written in a more general form as $\rho(t; \lambda) = \lambda^2 \rho(t/\lambda)$ with $\rho(0) = 0$ and $1 - t \leq (d/dt)\rho(t) \leq 1$ for $t > 0$, including the SCAD with $(d/dt)\rho(t) = 1 \wedge (1 - (t - 1)/(\gamma - 1))_+$, $\gamma \geq 2$, and the MCP with $(d/dt)\rho(t) = 1 \wedge (1 - t/\gamma)_+$, $\gamma \geq 1$. It can be verified that the ℓ_α penalty for $0 \leq \alpha \leq 1$, the SCAD and MCP are all concave. Another simple concave penalty is $\rho(t; \lambda) = \min(\lambda^2 \gamma / 2, \lambda |t|)$, $\gamma \geq 1$, introduced in [45] as capped- ℓ_1 penalty.

The above mentioned nonconvex interpolations of ℓ_0 and ℓ_1 penalties typically gain smoothness over the ℓ_0 penalty and thus allow more computational options. Meanwhile, they may improve variable selection accuracy and gain oracle properties by reducing the bias of Lasso. A more direct way to reduce the bias of Lasso is via the adaptive Lasso procedure [49], which solves the following weighted ℓ_1 regularization problem for some $\alpha \in (0, 1)$:

$$\min_{\mathbf{b} \in \mathbb{R}^p} \left[\frac{1}{2n} \|\mathbf{y} - \mathbf{X}\mathbf{b}\|_2^2 + \lambda \sum_{j=1}^p |\hat{w}_j|^{-\alpha} |b_j| \right],$$

where \hat{w} is an estimator of β (for example, the solution of the standard unweighted Lasso with regularization parameter λ). A low-dimensional analysis in [49] showed that the Adaptive Lasso solution can achieve the oracle property asymptotically. A high dimensional analysis of this procedure was given in [18]. For variable selection consistency and oracle properties to hold, the adaptive Lasso requires stronger conditions in terms of the minimum signal strength $\min_{j \in \text{supp}(\beta)} |\beta_j|$ than what is optimal. Specifically, the optimal requirement is $\min_{j \in \text{supp}(\beta)} |\beta_j| \geq \gamma \lambda_{univ}$ with $\lambda_{univ} = \sigma \sqrt{(2/n) \ln p}$ for some constant γ that may depend on an ℓ_2 regularity condition (also see Eq (11) below), which can be achieved by other procedures [42, 47]; however, adaptive Lasso requires $\min_{j \in \text{supp}(\beta)} |\beta_j|$ to be significantly larger than the optimal order of λ_{univ} . This means adaptive Lasso is sub-optimal for sparse estimation problems. We also observe that adaptive Lasso does not directly minimize a concave loss function, and hence it is not an instance of (2). It was later noted that this procedure is only one iteration of using the so-called MM (majorization-minimization) principle to solve (2) with bridge penalty (for example, see [50]). The corresponding MM procedure is referred to as multi-stage convex relaxation in [45, 47]. For sparse estimation problem (2) with a penalty $\rho(t; \lambda)$ that is concave in $|t|$, this method iteratively invokes the solution of the following reweighted ℓ_1 regularization problem for stage $\ell = 1, 2, \dots$, starting with the initial

value of $\widehat{\boldsymbol{\beta}}^{(0)} = 0$:

$$\widehat{\boldsymbol{\beta}}^{(\ell)} = \arg \min_{\mathbf{b} \in \mathbb{R}^p} \left[\frac{1}{2n} \|\mathbf{X}\mathbf{b} - \mathbf{y}\|_2^2 + \sum_{j=1}^p \lambda_j^{(\ell)} |b_j| \right],$$

where $\lambda_j^{(\ell)} = (\partial/\partial t)\rho(t; \lambda)|_{t=|\widehat{\beta}_j^{(\ell-1)}|}$ ($j = 1, \dots, p$). This procedure may be regarded as a multi-stage extension of adaptive Lasso, which corresponds to the stage-2 solution $\widehat{\boldsymbol{\beta}}^{(2)}$ with bridge penalty. Unlike results for adaptive Lasso, the results in [45, 47] for the multistage relaxation method allow $\min_{j \in \text{supp}(\boldsymbol{\beta})} |\beta_j|$ to achieve the optimal order of λ_{univ} , which match those of [42] and improve upon [18]. Moreover, only $\ell = O(\ln(\|\boldsymbol{\beta}\|_0))$ stages is necessary in order to achieve model selection consistency and oracle properties. It is worth pointing out that the multi-stage procedure can also be adapted to work with the Dantzig selector formulation [23].

For large p , the global solution of a nonconvex regularization method is hard to compute, so that local solutions are often used instead. Therefore theoretical analysis of nonconvex regularization has so far focused on specific numerical procedures that can find local solutions. For the ℓ_0 penalty, the penalized loss in (2) is typically evaluated for a subset of the 2^p possible models $\text{supp}(\mathbf{b})$ such as those generated in stepwise regression. For smooth concave penalties, iterative algorithms can be used to find local minima of the penalized loss in (2) for a set of penalty levels [19, 50, 45, 4, 25, 47]. For the MCP and other quadratic spline concave penalties, a path following algorithm can be used to find local minima for an interval of penalty levels [42].

Advances have been carried out in the analysis of nonconvex regularization methods in multiple fronts [14, 15, 49, 18, 46, 42, 45, 5]. For concave penalized loss in (2), local minimizers exist with the oracle property (3) under mild conditions [14, 15]. However, it remains unclear whether there exist computationally efficient procedures that can find local minimizers investigated in [14, 15]. For the MCP, the local minima generated by the path following algorithm controls the estimation error and model size in the sense of (8) under an ℓ_2 regularity condition on \mathbf{X} [42]. Under the additional condition

$$\min_{\beta_j \neq 0} |\widehat{\beta}_j^o| \geq \gamma \lambda_{univ} \geq \sup \{t : (\partial/\partial t)\rho(t; \lambda) \neq 0\} \quad (11)$$

with $\lambda_{univ} = \sigma\sqrt{(2/n)\ln p}$ and a certain constant $\gamma > 1$, the same path following solution has the oracle property (3) and thus the sign-consistency property [42]. Similar results hold for the SCAD and certain other quadratic spline penalties [42]. Under (11) and ℓ_2 regularity conditions on \mathbf{X} , the oracle property (3) and model selection consistency has also been established for a specific forward/backward stepwise regression scheme [46] that can be regarded as an approximate ℓ_0 penalty minimization algorithm. As we have mentioned earlier, the multi-stage relaxation scheme for minimizing (2) also leads to oracle inequality and model selection consistency under (11) and ℓ_2 regularity conditions on \mathbf{X} [45, 47].

While a number of specialized results were obtained for specialized numerical procedures under appropriate conditions, it is not clear what are the relationship among these solutions. For example, it is not clear whether the global solution of (2) is unique and whether it corresponds to solutions of various numerical procedures studied in the literature. This leads to a conceptual gap in the sense that it is not clear whether we should study specific local solutions as in the above mentioned previous work or we should try to solve (2) as accurately as possible (with the hope of finding the global solution). It is worth mentioning that related to this question, oracle inequalities involving

global solutions with nonconvex penalties have been studied in the literature (for example, see related sections in [5]). However, such oracle inequalities do not lead to results comparable to those of [42, 45, 47]. Another relevant study is [20], which showed that in the lower dimensional scenario with $p \leq n$, the global solution of (2) agrees with the oracle estimator $\hat{\beta}^o$ for the SCAD penalty when $\min_{\beta_j \neq 0} |\hat{\beta}_j^o|$ is sufficiently large, and some other appropriate assumptions hold. However, their analysis does not directly generalize to the more complex high dimensional setting.

The purpose of the remaining of this paper is to present some general results showing that under appropriate ℓ_2 -regularity conditions, the global solution of an appropriate nonconvex regularization method leads to desirable recovery performance; moreover, under suitable conditions, the global solution corresponds to the unique sparse local solution, which can be obtained via different numerical procedures. This leads to a unified view of concave high dimensional sparse estimation methods that can serve as a guideline to develop additional numerical algorithms for concave regularization.

3 High-Level Description of Main Results

As we have discussed in our brief survey, concave regularized methods have been proven to control the estimation error and the dimension of the selected model (8) under ℓ_2 regularity conditions and possess the oracle property (3) or the sign-consistency property under the additional assumption (11). However, these results are established for specific local solutions of (2) with specific penalties. For $p > n$ it is still unclear if the global minimizer in (2) is identical to these local solutions or controls estimation and selection errors in a similar way. In this paper, we unify the aforementioned results with the global solution of (2). Technical results are rigorously described in Section 4 below. This section explains the main thrust of these results.

We are mainly interested in two situations: ℓ_0 regularization where $\rho(t; \lambda)$ is discontinuous at $t = 0$, and smooth regularization which is continuous for all $t \geq 0$ and piece-wise differentiable. However, our basic results require only sub-additivity and monotonicity of $\rho(t; \lambda)$ in t in $[0, \infty)$.

We shall first describe assumptions of our analysis in Subsection 4.1. As we have pointed out, the key regularity conditions required in our analysis are expressed in terms of the sparse eigenvalues in (4) or invertibility factors RIF and CIF defined in (14) and (5). For the sake of clarity, we assume that these quantities are all constants, and this requirement is an ℓ_2 regularity condition. Another condition required by our analysis is called *null-consistency*, which requires that if $\beta = 0$, then the global minimizer of (2) is achievable at $\hat{\beta} = 0$ (the actual condition, given in Assumption 2, is slightly stronger). Clearly this condition depends both on the matrix \mathbf{X} and on the noise vector ε . It is shown in Subsection 4.1 that under the standard sub-Gaussian noise assumption (see Assumption 1), the null-consistency condition is ℓ_2 regular.

In summary, all assumptions on \mathbf{X} needed in our analysis are ℓ_2 regular; with this in mind, we may examine the main results, which are divided into four subsections.

Subsection 4.2 is concerned with basic properties of global optimal solution of (2) for all subadditive nondecreasing penalties. Theorem 1 gives ℓ_q -norm error bounds for $\|\hat{\beta} - \beta\|_q$ and a bound of the prediction error $\|\mathbf{X}\hat{\beta} - \mathbf{X}\beta\|_2$ that are comparable with known results for ℓ_1 regularization. This means that under appropriate ℓ_2 regularity conditions, the global solution of concave regularization problems are no worse than the Lasso solution in terms of the order of estimation error. Theorem 2 shows that the global optimal solution of (2) is sparse, and under appropriate ℓ_2 regularity conditions, the sparsity is of the same order as $\|\beta\|_0$; that is, $\|\hat{\beta}\|_0 = O(\|\beta\|_0)$. Thus, (8) holds for the global solution of (2). Moreover, if the second order derivative of $\rho(t; \lambda)$ with

respect to t is sufficiently small, then the global solution is also the unique sparse local solution of (2). That is if a vector $\tilde{\beta}$ is a local solution of (2) which is sparse: $\|\tilde{\beta}\|_0 = O(\|\beta\|_0)$, then $\tilde{\beta}$ is the global solution of (2). None of these results require that $\min_{\beta_j \neq 0} |\hat{\beta}_j^o|$ to be bounded away from zero. Furthermore, since these results require only ℓ_2 regularity conditions, they apply to the case of $p \gg n$ as long as $s^*(\ln p)/n$ is small.

Subsection 4.3 contains results specifically for ℓ_0 regularization. First, the global solution of ℓ_0 regularization is sparse. Moreover, with sub-Gaussian noise, the prediction error bound for ℓ_0 penalty in Theorem 3 does not depend on properties of the design matrix \mathbf{X} . This significantly improves upon the corresponding result for general penalties in Theorem 1, which requires a non-trivial RIF₁ condition on the design matrix \mathbf{X} . If the smallest sparse eigenvalue of $\mathbf{X}^\top \mathbf{X}/n$ is bounded from below, then we obtain in Theorem 4 the selection consistency for ℓ_0 regularization under (11), which implies the oracle property.

Subsection 4.4 considers penalties $\rho(t; \lambda)$ which are both left- and right-differentiable, for which one can define (approximate) local solutions that are what numerical optimization procedures compute. Theorem 5 considers the distance between two approximate local solutions. An immediate consequence of the result says that under appropriate assumptions, if $(\partial/\partial t)\rho(t; \lambda) = 0$ when t is sufficiently large, then there is a unique sparse local solution of (2) that corresponds to the oracle least squares solution $\hat{\beta}^o$ under (11). Therefore the unique local solution has the oracle property. Moreover, this unique local solution has to be the global optimal solution according to Theorem 2. While Theorem 5 shows that it is possible for a penalty that is not second order differentiable to have a unique sparse local solution, it requires the condition (11) for such penalties. In contrast, with a second order differentiable concave penalty, condition (11) is not needed in Theorem 5 for sparse local solutions to be unique. This suggests an advantage for using smooth concave penalties which may lead to fewer local solutions under certain conditions. Theorem 6 gives sufficient conditions under which the global optimal solution of (2) achieves model selection consistency. These sufficient conditions generalize the irrepresentable condition (9) for the model selection consistency of Lasso. However, unlike the irrepresentable condition for Lasso, which is not an ℓ_2 regularity condition, for a concave penalty where $(\partial/\partial t)\rho(t; \lambda)$ is small for sufficiently large t , the generalized irrepresentable condition required in Theorem 6 automatically holds when $\min_{\beta_j \neq 0} |\hat{\beta}_j^o|$ is not too small. Moreover, for appropriate nonconvex penalties, it is possible to achieve a selection threshold of optimal order as in (11).

Note that results in Subsection 4.4 show that if one can find a local solution of (2) and the solution is sparse, then under appropriate conditions, it is the global solution of (2) and it is close to the oracle least squares solution $\hat{\beta}^o$. It is possible to design numerical procedures that find a sparse local solution of (2). For such a procedure, results of Subsection 4.4 directly applies. Subsection 4.5 further develops along this line of thinking. Theorem 7 shows that if a local solution is also an approximate global solution, then it is sparse. This fact can be combined with results in Subsection 4.4 to imply that under appropriate conditions, this particular local solution is the unique sparse local solution (which is also the global solution). Moreover, such a solution can be obtained via Lasso followed by gradient descent, as it can be shown that Lasso is a sufficiently accurate approximate global solution of (2) for the result to apply.

Our results essentially imply the following: under appropriate ℓ_2 regularity conditions, plus appropriate assumptions on the penalty $\rho(t; \lambda)$, procedures considered earlier such as MCP [42] or multi-stage convex relaxation [19, 50, 45] give the same local solution that is also the global minimizer of (2). Moreover, other procedures (such as Lasso followed by gradient descent) can be

designed to obtain the same solution. Therefore these results present a coherent view of concave regularization by unifying a number of earlier approaches and by extending a number of previous results. This unified theory presents a more satisfactory treatment of concave high dimensional sparse estimation procedures.

4 Technical Statements of the Main Results

This section describes in detail our new technical results characterizing the global and local optimal solutions of (2) under different regularization conditions. Before going into the main results, we will specify some assumptions and definitions required in our analysis.

4.1 General Assumptions and Definitions

In this subsection, we describe and discuss general conditions imposed in the rest of the paper.

We first consider conditions on the regularizer $\rho(t; \lambda)$. We assume throughout the sequel the following conditions on the penalty function:

- (i) $\rho(0; \lambda) = 0$;
- (ii) $\rho(-t; \lambda) = \rho(t; \lambda)$;
- (iii) $\rho(t; \lambda)$ is non-decreasing in t in $[0, \infty)$;
- (iv) $\rho(t; \lambda)$ is subadditive with respect to t , $\rho(x + y; \lambda) \leq \rho(x; \lambda) + \rho(y; \lambda)$ for all $x, y \geq 0$.

This family of penalties is closed under the summation and maximization operations and includes all functions increasing and concave in $|t|$. Although we are mainly interested in the case where $\rho(t; \lambda)$ is concave in $|t|$, all of our results hold under the above specified weaker conditions, sometimes with side conditions such as the monotonicity of $\rho(t; \lambda)/t$ for $t > 0$ and the continuity of $\rho(t; \lambda)$ at $t = 0$. Therefore we will mention explicitly when such side conditions are needed.

We are particularly interested in the ℓ_0 regularization $\rho(t; \lambda) = (\lambda^2/2)I(t \neq 0)$ which is discontinuous at $t = 0$. In addition, we are interested in regularizer $\rho(t; \lambda)$ that is continuous in $t \geq 0$ and piece-wise differentiable. With such regularizers, local solutions of (2) can be defined as solutions with gradient zero. A local solution can be obtained using standard numerical procedures such as gradient descent.

Given a regularizer $\rho(t; \lambda)$ and any fixed $\lambda > 0$, we define the threshold level of the penalty as

$$\lambda^* := \inf_{t>0} \{t/2 + \rho(t; \lambda)/t\}. \quad (12)$$

The quantity λ^* is a function of λ that provides a natural normalization of λ . We call λ^* the threshold level since $\arg \min_t \{(z - t)^2/2 + \rho(t; \lambda)\} = 0$ iff $|z| \leq \lambda^*$. This can be easily seen from $(z - t)^2/2 + \rho(t; \lambda) - z^2/2 = t\{t/2 + \rho(t; \lambda)/t - z\}$. If $\rho(t; \lambda)$ is continuous at $t = 0$ and concave in $t \in (0, \infty)$, then $\lambda^* \leq \lim_{t \rightarrow 0+} (\partial/\partial t)\rho(t; \lambda)$. For simplicity, we may also require that $\rho(t; \lambda)$ be chosen such that $\lambda^* = \lambda$, which holds for the penalties discussed in Subsection 2.2, such as ℓ_0 , bridge, SCAD, MCP, and capped- ℓ_1 penalties.

In the following and in the proofs, we will use the short-hand notation

$$\|\rho(\mathbf{b}; \lambda)\|_1 = \sum_{j=1}^p \rho(b_j; \lambda), \quad \forall \mathbf{b} = (b_1, \dots, b_p)^\top.$$

Definition 2. The following quantity bounds a general penalty via ℓ_1 penalty for sparse vectors:

$$\Delta(a, k; \lambda) = \sup \left\{ \|\rho(\mathbf{b}; \lambda)\|_1 : \|\mathbf{b}\|_1 \leq ak, \|\mathbf{b}\|_0 = k \right\}. \quad (13)$$

Proposition 1. Let $\rho^*(t; \zeta) = \zeta|t| + (\zeta - |t|/2)_+^2/2$. Let λ^* be as in (12). Then,

$$\begin{aligned} \min \left\{ \lambda^*|t|/2, (\lambda^*)^2/2 \right\} &\leq \rho(t; \lambda) \leq \rho^*(t; \lambda^*), \\ \Delta(a, k; \lambda) &\leq k\rho^*(a; \lambda^*) \leq k \max(a, 2\lambda^*)\lambda^*. \end{aligned}$$

Remark 3. It follows from Proposition 1 that given a threshold level λ^* , all penalty functions satisfying general conditions (i)-(iv) are bounded by a capped- ℓ_1 penalty from below and the maximum of the ℓ_0 and ℓ_1 penalties from above, up to a factor of 2. The function $\rho^*(t; \zeta)$ is a convex quadratic spline fit of $\max(\zeta^2/2, \zeta|t|)$, the maximum of the ℓ_0 and ℓ_1 penalties with threshold level ζ .

Remark 4. A trivial upper bound is $\Delta(a, k; \lambda) \leq k \max_t \rho(t; \lambda)$, which is useful only for bounded penalties. The ℓ_∞ bound $\rho(t; \lambda) \leq \gamma^* \lambda^2$ holds with $\gamma^* = 1/2$ for the ℓ_0 penalty, $\gamma^* = \gamma/2$ for the capped- ℓ_1 penalty and MCP, and $\gamma^* = (1 + \gamma)/2$ for the SCAD penalty. If $\rho(t; \lambda)$ is concave in $t \in [0, \infty)$, then $\Delta(a, k; \lambda) \leq k\rho(a; \lambda)$ by the Jensen inequality. For $a \geq 2\lambda^*$, $\Delta(a, k; \lambda) \leq a\lambda^*k$ matches the trivial bound for the ℓ_1 penalty, for which $\lambda = \lambda^*$.

Next, we consider conditions on the design matrix \mathbf{X} . Recall that \mathbf{X} is column normalized to $\|\mathbf{x}_j\|_2^2 = n$ for simplicity. Our analysis also depends on the sparse eigenvalues defined in (4) and the restricted invertibility factor defined as follows.

Definition 3. For $q \geq 1$, $\xi > 0$ and $S \subset \{1, \dots, p\}$, we define the restricted invertibility factor as

$$\text{RIF}_q(\xi, S) = \inf \left\{ \frac{|S|^{1/q} \|\mathbf{X}^\top \mathbf{X} \mathbf{u}\|_\infty}{n \|\mathbf{u}\|_q} : \|\rho(\mathbf{u}_{S^c}; \lambda)\|_1 < \xi \|\rho(\mathbf{u}_S; \lambda)\|_1 \right\}. \quad (14)$$

The restricted invertibility factor is the quantity needed to separate conditions on \mathbf{X} and $\boldsymbol{\varepsilon}$ in our analysis. For $1 \leq q \leq 2$, sparse eigenvalues can be used to find lower bounds of $\text{RIF}_q(\xi, S)$.

Proposition 2. Let CIF be as in (5). If $t/\rho(t; \lambda)$ is increasing in $t \in (0, \infty)$, then

$$\text{RIF}_q(\xi, S) \geq \inf_{|A|=|S|} \text{CIF}_q(\xi, A). \quad (15)$$

Remark 5. For the ℓ_1 penalty, $\text{RIF}_q = \text{CIF}_q$. If $\rho(t; \lambda)$ is concave in $t \in [0, \infty)$, then $t/\rho(t; \lambda)$ is increasing in t . Thus, Proposition 2 is applicable to all penalty functions discussed in Subsection 2.2, including the ℓ_0 , bridge, SCAD, MCP, and capped- ℓ_1 penalties.

Remark 6. The CIF can be uniformly bounded from below in terms of sparse eigenvalues:

$$\text{CIF}_q(\xi, S) \geq \frac{I\{1 \leq q \leq 2\} \{\kappa_-(k + \ell) - (\xi/2)(k/\ell)^{1/2} \kappa_+(k + 5\ell)\}}{(1 + \xi)^{2/q-1} (1 + \xi^2 k/(4\ell))^{1-1/q} (1 + \ell/k)^{1/2}}, \quad (16)$$

for all $1 \leq \ell \leq (p - |S|)/5$ by Proposition 5 and (21) in [41], where $k = |S|$, and $\kappa_-(m)$ and $\kappa_+(m)$ are as in (4). For example, if we take $\xi = 2$ and $\ell = 2k$ and $q = 2$, then

$$\text{CIF}_2(\xi, S) \geq \{\kappa_-(3k) - \kappa_+(11k)/\sqrt{2}\}/\sqrt{4.5}.$$

Remark 7. It follows from Proposition 2 and Remark 6 that conditions $\text{RIF}_q(\xi, S) > 0$ and $1/\text{RIF}_q(\xi, S) = O(1)$ are both ℓ_2 -regularity conditions on \mathbf{X} for $1 \leq q \leq 2$. Moreover, $\text{rank}(\mathbf{X}) = p$ implies $\text{RIF}(\xi, S) > 0$. To check the ℓ_2 regularity of these conditions, we suppose that the rows of \mathbf{X} are iid from $N(0, \Sigma)$ with all eigenvalues of Σ in $[c_1, c_2] \subset (0, \infty)$. Then, $c_1/2 \leq \kappa_-(m)$ and $\kappa_+(m) \leq 2c_2$ with at least probability $1 - \delta \in [0, 1)$ for $m \leq c_3 n / \ln(p/\delta)$ for a certain $c_3 > 0$. Let $c_4 = \{c_1/(\xi c_2)\}^2$. In this event, setting $k = s^*$ and $\ell = (m - s^*)/5$ in (16) yields

$$\min_{|S| \leq s^*} \text{RIF}_2(\xi, S) \geq \min_{|S| \leq s^*} \text{CIF}_2(\xi, S) \geq (c_1/4) / \sqrt{(1 + \xi^2 c_4/4)(1 + 1/c_4)}$$

when $5s^*/(m - s^*) < c_4$ for some $m \leq c_3 n / \ln(p/\delta)$, which holds when $(s^*/n) \ln(p/\delta) \leq c_3/(1+5/c_4)$.

Finally, we consider conditions on the error vector.

Assumption 1. An error vector $\boldsymbol{\varepsilon}$ is sub-Gaussian with noise level σ if for all $t \geq 0$:

$$P(|\mathbf{u}^\top \boldsymbol{\varepsilon}| > \sigma t) \leq \exp(-t^2/2)$$

for all vector \mathbf{u} with $\|\mathbf{u}\|_2 = 1$ and

$$P(\|\mathbf{P}_A \boldsymbol{\varepsilon}\|_2 / |A|^{1/2} > \sigma(1+t)) \leq \exp(-|A|t^2/2)$$

for all subsets $A \subset \{1, \dots, p\}$, where \mathbf{P}_A is the orthogonal projection to the range of \mathbf{X}_A (that is, $\mathbf{P}_A = \mathbf{X}_A \mathbf{X}_A^\dagger$, where \mathbf{X}_A^\dagger is the Moore-Penrose generalized inverse of \mathbf{X}_A).

The above sub-Gaussian condition holds with $\boldsymbol{\varepsilon} \sim N(0, \sigma^2 \mathbf{I}_{n \times n})$. It is equivalent to the more common version of the sub-Gaussian condition $E e^{\mathbf{v}^\top \boldsymbol{\varepsilon} / \sigma'} \leq e^{\|\mathbf{v}\|^2/2}$ for all vectors \mathbf{v} and a constant σ' of the same order as σ . As we have mentioned in Section 3, what we really need is a null-consistency condition, which we give below. The sub-Gaussian condition will be used to verify the null consistency condition.

Assumption 2. Let $\eta \in (0, 1]$. We say that the regularization method (2) satisfies the η null-consistency condition if the following equality holds:

$$\min_{\mathbf{b} \in \mathbb{R}^p} \left(\|\boldsymbol{\varepsilon}/\eta - \mathbf{X}\mathbf{b}\|_2^2 / (2n) + \|\rho(\mathbf{b}; \lambda)\|_1 \right) = \|\boldsymbol{\varepsilon}/\eta\|_2^2 / (2n). \quad (17)$$

Remark 8. Given $\eta = 1$, the null-consistency condition means that if $\boldsymbol{\beta} = 0$, then the global minimizer of (2) is achievable at $\hat{\boldsymbol{\beta}} = 0$. This requirement is clearly necessary for the global minimizer of (2) to satisfy the error bound (19) in Theorem 1 below for $|S| = 0$. Here, we also allow a slightly stronger condition with $\eta < 1$, which requires $\hat{\boldsymbol{\beta}} = 0$ for $\boldsymbol{\beta} = 0$ when the noise $\boldsymbol{\varepsilon}$ is proportionally inflated by $1/\eta$.

Proposition 3. Suppose that $\boldsymbol{\varepsilon}$ is sub-Gaussian with noise level σ , $0 < \delta \leq 1$ and $\zeta_0 > 0$. Suppose $\rho(t; \lambda) \geq ((\lambda^*)^2/2) \wedge (\lambda^*|t|)$ with $\lambda^* \geq (1 + \zeta_0)(\sigma/\eta)n^{-1/2}(1 + \sqrt{2 \ln(2p/\delta)})$. Then, (2) satisfies the η null-consistency condition with at least probability $2 - e^{\delta/2} - \exp(-n(1 - 1/\sqrt{2})^2)$, provided that

$$\max \left\{ \lambda_{\max}^{1/2}(\mathbf{X}_B^\top \mathbf{P}_A \mathbf{X}_B / n) : \begin{array}{l} B \cap A = \emptyset, |A| = \text{rank}(\mathbf{P}_A) = |B| = k, \\ k(1 + \zeta_0)^2(1 + \sqrt{2 \ln(2p/\delta)})^2 \leq 2n \end{array} \right\} \leq \zeta_0. \quad (18)$$

Moreover, (18) holds with no smaller probability than $1 - \delta^4/(16p^2)$ if the rows of \mathbf{X} are iid from $N(0, \Sigma)$ and $\sqrt{8} \lambda_{\max}^{1/2}(\Sigma) \leq \zeta_0(1 + \zeta_0)$. This means that under the sub-Gaussian condition on $\boldsymbol{\varepsilon}$, the η null-consistency is an ℓ_2 -regularity condition.

Remark 9. The condition $\rho(t; \lambda) \geq \min(\lambda^2/2, \lambda|t|)$ holds for the ℓ_0 , ℓ_1 , SCAD, and capped ℓ_1 penalties, so that Proposition 3 is directly applicable with $\lambda = \lambda^*$. In general, the condition of Proposition 3 holds for all penalties considered in this paper when the threshold level in (12) satisfies $\lambda^* \geq 2(1 + \zeta_0)(\sigma/\eta)n^{-1/2}(1 + \sqrt{2\ln(2p/\delta)})$, in view of the lower bound of $\rho(t; \lambda)$ in Proposition 1. For ℓ_0 and ℓ_1 penalties, we may set $\zeta_0 = 0$ in Proposition 3 (the extra condition (18) is not necessary). The simplified condition for ℓ_0 penalty is explicitly given in Theorem 3. For the ℓ_1 penalty, the η null consistency condition is equivalent to $\|\mathbf{X}^\top \boldsymbol{\varepsilon}\|_\infty \leq \eta\lambda n$.

4.2 Basic Properties of the Global Solution

We now turn our attention to the global solution of (2) with a general subadditive nondecreasing regularizer $\rho(t; \lambda)$. We first consider the estimation of $\mathbf{X}\boldsymbol{\beta}$ and $\boldsymbol{\beta}$.

Theorem 1. Let $S = \text{supp}(\boldsymbol{\beta})$, $\widehat{\boldsymbol{\beta}}$ be as in (2), λ^* as in (12), and $\text{RIF}_q(\xi, S)$ as in (14). Consider $\eta \in (0, 1)$, and $\xi = (\eta + 1)/(1 - \eta)$, and assume that (17) holds. Then for all $q \geq 1$:

$$\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_q \leq (1 + \eta)\lambda^*|S|^{1/q}/\text{RIF}_q(\xi, S), \quad (19)$$

and with $a_1 = (1 + \eta)/\text{RIF}_1(\xi, S)$ and $\Delta(a, k; \lambda)$ in (13),

$$\|\mathbf{X}\widehat{\boldsymbol{\beta}} - \mathbf{X}\boldsymbol{\beta}\|_2^2/n \leq 2\xi\Delta(a_1\lambda^*, |S|; \lambda) \leq 2\xi(a_1 \vee 2)(\lambda^*)^2|S|. \quad (20)$$

By using the bound $\Delta(a_1\lambda^*, |S|; \lambda) \leq |S|\max_t \rho(t; \lambda)$, we obtain the following corollary.

Corollary 1. Consider penalties $\rho(t; \lambda)$ indexed by the threshold level; $\lambda^* = \lambda$ in (12). Suppose that the η null consistency condition (17) holds. Let $S = \text{supp}(\boldsymbol{\beta})$ and $\gamma^* = \max_t \rho(t; \lambda)/\lambda^2$. Then,

$$\|\mathbf{X}\widehat{\boldsymbol{\beta}} - \mathbf{X}\boldsymbol{\beta}\|_2^2/n \leq 2\{(1 + \eta)/(1 - \eta)\}\gamma^*\lambda^2|S|.$$

In particular, $\gamma^* = 1/2$ for the ℓ_0 penalty $\rho(t; \lambda) = (\lambda^2/2)I(t \neq 0)$, $\gamma^* = \gamma/2$ for the capped- ℓ_1 penalty $\rho(t; \lambda) = (\lambda^2\gamma/2) \wedge (\lambda|t|)$ and the MCP $\rho(t; \lambda) = \lambda \int_0^{|t|} (1 - x/(\lambda\gamma))_+ dx$, and $\gamma^* = (1 + \gamma)/2$ for the SCAD penalty $\rho(t; \lambda) = \lambda \int_0^{|t|} \min\{1, (1 - (x/\lambda - 1)/(\gamma - 1))_+\} dx$.

Remark 10. It is worthwhile to note that the prediction error bound in Corollary 1 does not depend on \mathbf{X} , provided that penalty is large enough to guarantee null consistency. For the ℓ_0 penalty, the null consistency requires only $\|\mathbf{x}_j\|_2 = \sqrt{n}$ on \mathbf{X} , which we assume anyway. For other concave penalties in Corollary 1, we are only able to provide null consistency in Proposition 3 under a mild condition on the upper eigenvalue of $\mathbf{X}_B^\top \mathbf{P}_A \mathbf{X}_B/n$, but not on the sparse lower eigenvalue of the Gram matrix.

Next we provide an upper bound for the sparseness of $\widehat{\boldsymbol{\beta}}$ based on Theorem 1 and the maximum sparse eigenvalue $\kappa_+(m)$. We denote by $\dot{\rho}(t; \lambda) = (\partial/\partial t)\rho(t; \lambda)$ any value between the left- and right- derivatives of $\rho(\cdot; \lambda)$ and assume the left- and right-differentiability of $\rho(\cdot; \lambda)$ whenever the notation $\dot{\rho}(t; \lambda)$ is invoked. For example, if $\rho(t; \lambda) = \lambda|t|$, then $\dot{\rho}(0\pm; \lambda) = \pm\lambda$ and $\dot{\rho}(0; \lambda)$ can be any value in $[-\lambda, \lambda]$ (which in all of our results, can be chosen as the most favorable value unless explicitly mentioned otherwise).

Theorem 2. Let $\{S, \hat{\beta}, \lambda^*, \eta, \xi, a_1\}$ and $\Delta(a, k; \lambda)$ be as in Theorem 1, and $\hat{S} = \text{supp}(\hat{\beta})$. Suppose that (17) holds. Consider $t_0 \geq 0$ and integer $m_0 \geq 0$ satisfying $m_0 = 0$ for $t_0 = 0$ and

$$\sqrt{2\xi\kappa_+(m_0)\Delta(a_1\lambda^*, |S|; \lambda)/m_0} + \|\mathbf{X}^\top \varepsilon/n\|_\infty < \inf_{0 < s < t_0} \dot{\rho}(s; \lambda) \quad (21)$$

for $t_0 > 0$. Then,

$$|\hat{S} \setminus S| < m := m_0 + \lfloor \xi \Delta(a_1\lambda^*, |S|; \lambda) / \rho(t_0; \lambda) \rfloor. \quad (22)$$

The η null consistency implies $\|\mathbf{X}^\top \varepsilon/n\|_\infty \leq \eta\lambda^*$ by Lemma 1 in Section 5. If $\rho(t; \lambda)$ is concave in $t > 0$, then the right-hand side of (21) can be replaced by $\dot{\rho}(t_0; \lambda)$ and $\rho(t_0; \lambda) \geq t_0\dot{\rho}(t_0; \lambda)$. These facts give the following corollary for ℓ_∞ bounded and ℓ_1 penalties.

Corollary 2. (i) Let $\rho(t; \lambda)$ and γ^* be as in Corollary 1. Suppose (2) is η null consistent in the sense of (17) and $\dot{\rho}(a_0\lambda; \lambda) \geq \lambda(1 - a_1/\gamma)$ for some $a_0 > 0$ and $a_1 \geq 0$. If $m_0 = \alpha|S|$ is an integer and $2\gamma^*\kappa_+(\alpha|S|)/\alpha < (1 - a_1/\gamma - \eta)^2(1 - \eta)/(1 + \eta)$, then

$$|\hat{S} \setminus S| < m := \left(\alpha + \frac{\gamma^*/a_0}{1 - a_1/\gamma} \right) |S|. \quad (23)$$

(ii) Let $\hat{S}^{(\ell_1)} = \text{supp}(\hat{\beta}^{(\ell_1)})$ with the Lasso (7) and CIF_q as in (5). In the event $\|\mathbf{X}^\top \varepsilon/n\|_\infty \leq \eta\lambda$,

$$\frac{2\kappa_+(\alpha|S|)/\alpha}{\text{CIF}_1((1 + \eta)/(1 - \eta), S)} < \frac{(1 - \eta)^3}{(1 + \eta)^2} \Rightarrow |\hat{S}^{(\ell_1)} \setminus S| < m := \alpha|S|. \quad (24)$$

Remark 11. Theorem 2 and Corollary 2 imply that the global solution $\hat{\beta}$ in (2) is sparse under appropriate assumptions. For ℓ_0 regularization, we may take $m_0 = t_0 = 0$ with the convention $\kappa_+(0)/0 = 0$ in (21). The Lasso also satisfies the dimension bound $|\hat{S} \setminus S| < m \vee 1$ under the SRC: $\{\kappa_+(m + |S|)/\kappa_-(m + |S|) - 1\}/(2 - 2a_0) \leq m/|S|$ with an $a_0 \in (0, 1)$, provided that $\lambda \geq (1 + o(1))\{\kappa_+^{1/2}(m)/a_0\}\sigma\sqrt{(2/n)\ln p}$ [42]. An advantage of (24) is to allow an λ not dependent on the upper sparse eigenvalue of the design for sub-Gaussian ε .

Remark 12. Let $\kappa^* = \sup_{0 < s < t} \{\dot{\rho}(t; \lambda) - \dot{\rho}(s; \lambda)\}/(t - s)$ be the maximum concavity of the penalty. Suppose $\kappa_- (|S| + m + \tilde{m} - 2) > \kappa^*$. Then, the penalized loss $L_\lambda(\mathbf{b})$ in (2) is convex in all models $\text{supp}(\mathbf{b}) = A$ with $|A \setminus S| \leq m + \tilde{m} - 2$. This condition has been called *sparse convexity* [42]. If m is as in (22) or (23) and $\hat{\beta}$ is a local solution of (2) with $\#\{j \notin S : \hat{\beta}_j \neq 0\} < \tilde{m}$, then the local solution must be identical to the global solution.

Remark 13. Consider penalties with $\lambda^* = \lambda$ which holds for all penalties discussed in Subsection 2.2. Let $\eta \in (0, 1)$ and $\lambda_* > 0$ be fixed. Suppose Theorem 2 or Corollary 2 is applicable with $m \leq \alpha^*|S|$ for a fixed constant α^* and all $\lambda \geq \lambda_*$. Suppose in addition $\dot{\rho}(t; \lambda)$ is continuous in $1/\lambda \in [0, 1/\lambda_*]$ uniformly in bounded sets of t . Under the sparse convexity condition $\kappa_- (|S| + m - 1) \geq \kappa_* > 0$, with the maximum concavity κ^* in Remark 12, the global solution forms a continuous path in \mathbb{R}^p as a function of $1/\lambda \geq 1/\lambda_*$. This path is identical to the output of the path following algorithm in [42] if it starts with $\hat{\beta} = 0$ at $1/\lambda = 0$. We will show in Theorem 7 that gradient algorithms beginning from the Lasso may also yield the global solution under the sparse convexity condition.

As a simple working example to illustrate Corollaries 1 and 2, we consider the capped- ℓ_1 penalty explicitly given in Corollary 1. Let $a_0 = \gamma/2$ in Corollary 2. We find

$$\begin{aligned} \|\mathbf{X}\hat{\boldsymbol{\beta}} - \mathbf{X}\boldsymbol{\beta}\|_2^2/n &\leq \lambda^2|S|\gamma(1+\eta)/(1-\eta), \\ \gamma\kappa_+(\alpha|S|) &\leq \alpha(1-\eta)^3/(1+\eta) \Rightarrow |\hat{S} \setminus S| < (\alpha+1)|S|. \end{aligned}$$

The MCP, also explicitly given in Corollary 1, provides the same prediction bound and

$$\gamma\kappa_+(\alpha|S|) \leq \alpha(2/3 - \eta)^2(1-\eta)/(1+\eta) \Rightarrow |\hat{S} \setminus S| < (\alpha + 9/4)|S|$$

by the same calculation with $\alpha_0 = \gamma/3$. Note that generally speaking, unless stronger conditions are imposed, Theorem 2 only implies that $|\hat{S} \setminus S| = O(|S|)$ but not $|\hat{S} \setminus S| = 0$ required for model selection consistency. The model selection consistency will be studied later in the paper.

4.3 The Global Solution of ℓ_0 Regularization

This subsection considers the global optimal solution $\hat{\boldsymbol{\beta}}^{(\ell_0)}$ of ℓ_0 regularization in (10). Our first result says that under appropriate conditions, this solution is sparse.

Theorem 3. *If for all $\mathbf{b} \in \mathbb{R}^p$: $\boldsymbol{\varepsilon}^\top \mathbf{X}\mathbf{b} \leq \lambda\eta\sqrt{n\|\mathbf{b}\|_0}\|\mathbf{X}\mathbf{b}\|_2$ for some $\eta < 1$, then (10) satisfies the η null-consistency condition. It implies that the global optimal solution of (10) satisfies*

$$\|\hat{\boldsymbol{\beta}}^{(\ell_0)}\|_0 \leq \frac{1+\eta^2}{1-\eta^2}\|\boldsymbol{\beta}\|_0, \quad \|\mathbf{X}\hat{\boldsymbol{\beta}}^{(\ell_0)} - \mathbf{X}\boldsymbol{\beta}\|_2^2 \leq \frac{(1+\eta)\lambda^2\|\boldsymbol{\beta}\|_0}{1-\eta}.$$

We also have the following result about model selection quality for ℓ_0 regularization.

Theorem 4. *Assume that the assumption of Theorem 3 holds. Let $s = 2\|\boldsymbol{\beta}\|_0/(1-\eta^2)$ and $\hat{\boldsymbol{\beta}}^o$ be as in (3). Suppose $\|\mathbf{X}^\top(\mathbf{P}_S\boldsymbol{\varepsilon} - \boldsymbol{\varepsilon})\|_\infty/n \leq \sqrt{2\kappa_-(s)}\lambda$, where \mathbf{P}_S is the orthogonal projection to the range of \mathbf{X}_S . Let $S = \text{supp}(\boldsymbol{\beta})$, $\delta^o = \#\{j \in S : |\hat{\beta}_j^o| < \lambda\sqrt{2/\kappa_-(s)}\}$, and $\hat{S} = \text{supp}(\hat{\boldsymbol{\beta}}^{(\ell_0)})$. Then,*

$$|S - \hat{S}| + 0.5|\hat{S} - S| \leq 2\delta^o, \quad \|\mathbf{X}(\hat{\boldsymbol{\beta}}^{(\ell_0)} - \hat{\boldsymbol{\beta}}^o)\|_2^2 \leq 2\lambda^2\delta^o.$$

If the error $\boldsymbol{\varepsilon}$ is sub-Gaussian in the sense of Assumption 1, then the condition of Theorems 3 and 4 holds with at least probability $2 - e^\delta$ for $\lambda \geq (\sigma/\eta)(1 + \sqrt{2\ln(p/\delta)})/\sqrt{n}$. Theorem 4 implies that model selection consistency can be achieved if the condition $\min_{j \in \text{supp}(\boldsymbol{\beta})} |\hat{\beta}_j^o| \geq \lambda/\sqrt{\kappa_-(s)}$ holds, which implies that $\delta^o = 0$.

4.4 Approximate Local Solutions

We have shown in Theorem 2 that under appropriate conditions, the global solution of (2) is sparse. If $\rho(t; \lambda)$ is both left- and right-differentiable, one can define the concept of local solution as follows. Given an excess $\nu \geq 0$, a vector $\tilde{\boldsymbol{\beta}} \in \mathbb{R}^p$ is an approximate local solution of (2) if

$$\|\mathbf{X}^\top(\mathbf{X}\tilde{\boldsymbol{\beta}} - \mathbf{y})/n + \dot{\rho}(\tilde{\boldsymbol{\beta}}; \lambda)\|_2^2 \leq \nu. \quad (25)$$

This $\tilde{\boldsymbol{\beta}}$ is a local solution if $\nu = 0$. Note that by convention, $\dot{\rho}(t; \lambda)$ can be chosen to be any value between $\dot{\rho}(t_-; \lambda)$ and $\dot{\rho}(t_+; \lambda)$ to satisfy the equation. In this subsection, we provide estimates of distances between approximate local solutions and use them to prove the equality of oracle

approximate local and global solutions of (2). This gives the selection consistency of the global solution studied in Subsection 4.2. The oracle LSE is considered as an approximate local solution. In addition, we define a sufficient condition for the existence of a sign consistent local solution which generalizes the irrepresentable condition for Lasso selection and becomes an ℓ_2 regularity condition on \mathbf{X} for a broad class of concave penalties.

We first provide estimates of distances between approximate local solutions. We use the following function $\theta(t, \kappa)$ to measure the degree of nonconvexity of a regularizer $\rho(t; \lambda)$ at $t \in \mathbb{R}$.

Definition 4. For $\kappa \geq 0$ and $t \in \mathbb{R}$, define

$$\theta(t, \kappa) := \sup_s \{-\text{sgn}(s - t)(\dot{\rho}(s; \lambda) - \dot{\rho}(t; \lambda)) - \kappa|s - t|\}.$$

Moreover, given $\mathbf{u} = (u_1, \dots, u_p)^\top \in \mathbb{R}^p$, we let $\theta(\mathbf{u}, \kappa) = [\theta(u_1, \kappa), \dots, \theta(u_p, \kappa)]$.

We are mostly interested in values of $\theta(t, \kappa)$ that achieves zero. We note that $\theta(t, \kappa) = 0$ for convex $\rho(t, \lambda)$ with $\kappa \geq 0$. More generally, let κ^* be the maximum concavity as in Remark 12. Then, $\theta(t, \kappa) = 0$ for all t iff $\kappa \geq \kappa^*$. For $\dot{\rho}(t+; \lambda) < \dot{\rho}(t-; \lambda)$, $\theta(t, \kappa) > 0$ for all finite κ . However, we only need $\theta(t, \kappa) = 0$ for a proper set of t in our selection consistency theory. As an example, for $\kappa = 2/\gamma$, the capped- ℓ_1 penalty $\rho(t; \lambda) = \min(\gamma\lambda^2/2, \lambda|t|)$ gives $\theta(t, \kappa) = 0$ when either $t = 0 \pm$ or $|t| \geq \gamma\lambda$.

The following theorem shows that under appropriate assumptions, two sparse approximate local solutions $\tilde{\beta}^{(1)}$ and $\tilde{\beta}^{(2)}$ are close.

Theorem 5. Let $\tilde{\beta}^{(j)}$ be approximate local solutions with excess $\nu^{(j)}$ and $\Delta = \tilde{\beta}^{(1)} - \tilde{\beta}^{(2)}$. Let $\kappa_\pm(\cdot)$ be the sparse eigenvalues in (4) and $\tilde{S}^{(j)} := \text{supp}(\tilde{\beta}^{(j)})$. Consider any $S \subset \{1, \dots, p\}$ with $k = |S|$, integer m such that $m + k \geq |\tilde{S}^{(1)} \cup \tilde{S}^{(2)}|$, and $0 < \kappa < \kappa_-(m + k)$. Then,

$$\|\mathbf{X}\Delta\|_2^2/n \leq \frac{2\kappa_-(m + k)}{(\kappa_-(m + k) - \kappa)^2} \left\{ \|\theta(|\tilde{\beta}_{\tilde{S}^{(1)}}^{(1)}|, \kappa)\|_2^2 + |\tilde{S}^{(2)} \setminus \tilde{S}^{(1)}| \theta^2(0+, \kappa) + \nu \right\} \quad (26)$$

with $\nu = \{(\nu^{(1)})^{1/2} + (\nu^{(2)})^{1/2}\}^2$, and

$$|S \setminus \tilde{S}^{(2)}| \leq \inf_{\lambda_0 > 0} \left[\#\left\{j \in S : |\tilde{\beta}_j^{(1)}| < \lambda_0 / \sqrt{\kappa_-(m + k)}\right\} + \|\mathbf{X}\Delta\|_2^2 / (\lambda_0^2 n) \right]. \quad (27)$$

If in addition $\theta(0+, \kappa) = 0$ and $\dot{\rho}(0+; \lambda) > \|\mathbf{X}_{S^c}^\top(\mathbf{X}\tilde{\beta}^{(1)} - \mathbf{y})/n\|_\infty$ with $S \supseteq S^{(1)}$ and $|S| \geq k$, then

$$|\tilde{S}^{(2)} \setminus S| \leq \frac{3[\{\kappa^2/\kappa_-(m + k) + \kappa_+(m)\} \|\mathbf{X}\Delta\|_2^2/n + \tilde{\nu}^{(2)}]}{\{\dot{\rho}(0+; \lambda) - \|\mathbf{X}_{S^c}^\top(\mathbf{X}\tilde{\beta}^{(1)} - \mathbf{y})/n\|_\infty\}^2}. \quad (28)$$

Let $S = \text{supp}(\beta)$. For comparison between a sparse local or global solution $\tilde{\beta}^{(2)}$ with $|\tilde{S}^{(2)} \setminus S| \leq m$ and an oracle solution $\tilde{\beta}^{(1)}$ with $\tilde{S}^{(1)} = S$, the sparse convexity condition implies $\tilde{\beta}^{(2)} = \tilde{\beta}^{(1)}$ when $\kappa^* < \kappa_-(|S| + m)$ as in Remark 12. However, since $\kappa^* = \infty$ when $\dot{\rho}(t+; \lambda) < \dot{\rho}(t-; \lambda)$ at a point $t > 0$, the sparse convexity argument requires the continuity of $\dot{\rho}(t; \lambda)$ for $t > 0$. This does not apply to the capped- ℓ_1 penalty. In Theorem 5, if $\theta(0+, \kappa) = \theta(\tilde{\beta}_S^{(1)}, \kappa) = 0$ with $\kappa < \kappa_-(|S| + m)$, then $\mathbf{X}\Delta = 0$, and hence $\tilde{\beta}^{(2)} = \tilde{\beta}^{(1)}$ (since $\kappa_-(|\tilde{S}^{(1)} \cup \tilde{S}^{(2)}|) > 0$). Thus, the sparse convexity

condition is much weakened to cover all left- and right-differentiable penalties such as the capped- ℓ_1 . On the other hand, Theorem 5 does not weaken the sparse convexity condition for the MCP $\rho(t; \lambda) = \lambda \int_0^{|t|} (1 - x/(\gamma\lambda))_+ dx$, for which $\theta(0+; \kappa) = 0$ iff $\kappa \geq \kappa^* = 1/\gamma$ iff $\theta(t; \kappa) = 0$ for all $t > 0$. It is worth pointing out that for a piecewise differentiable penalty that is not second order differentiable, the condition $\theta(\tilde{\beta}_S^{(1)}, \kappa) = 0$ (thus, the uniqueness of local solution) typically requires $|\tilde{\beta}_j^{(1)}|$ to be large to avoid the discontinuities of $\dot{\rho}(t; \lambda)$ when $j \in S$. As pointed out in Remark 12, this is not necessary when the penalty is second order differentiable. This means that there can be advantages of using smooth penalty terms that may have fewer local minimizers under certain conditions.

As a simple working example to illustrate Theorem 5, we consider the capped ℓ_1 penalty of the form $\rho(t; \lambda) = \min(\gamma\lambda^2/2, \lambda|t|)$. Let $S = \text{supp}(\beta)$. Assume that $\kappa = \kappa_-(m + |S|)/2 \geq 2/\gamma$. Then $\theta(t, \kappa) = 0$ when either $t = 0$ or $|t| \geq \gamma\lambda$. Therefore, if we define $\tilde{\beta}^{(1)}$ as $\tilde{\beta}_j^{(1)} = \hat{\beta}_j^o$ when $|\hat{\beta}_j^o| \geq \gamma\lambda$ and $\tilde{\beta}^{(j)} = 0$ otherwise, then

$$\|\mathbf{X}\Delta\|_2^2/n \leq \frac{8\nu}{\kappa_-(m + |S|)},$$

and by taking $\lambda_0 = \gamma\lambda\sqrt{\kappa_-(m + |S|)}$, we have

$$|S \setminus \tilde{S}^{(2)}| \leq \frac{\|\mathbf{X}\Delta\|_2^2}{\gamma^2\lambda^2\kappa_-(m + |S|)n}, \quad |\tilde{S}^{(2)} \setminus S| \leq \frac{3[1.25\kappa_+(m)\|\mathbf{X}\Delta\|_2^2/n + \tilde{\nu}^{(2)}]}{\{\lambda - \|\mathbf{X}_{S^c}^\top(\mathbf{X}\tilde{\beta}^{(1)} - \mathbf{y})/n\|_\infty\}^2}.$$

We now consider selection consistency of the global solution (2) by comparing it with an oracle solution with Theorem 5. For this purpose, we treat the oracle LSE as an approximate local solution by finding its excess ν in (25), and provide a sufficient condition for the existence of a sign consistent oracle local solution. This sufficient condition is characterized by the following extension of the quantities θ_1^* and θ_2^* in (9) from the ℓ_1 to general penalty:

$$\begin{aligned} \theta_1 &= \inf \{ \theta : \|(\mathbf{X}_S^\top \mathbf{X}_S/n)^{-1} \dot{\rho}(\mathbf{v}_S + \hat{\beta}_S^o; \lambda)\|_\infty \leq \theta\lambda^*, \forall \|\mathbf{v}_S\|_\infty \leq \theta\lambda^* \}, \\ \theta_2 &= \sup \{ \|\mathbf{X}_{S^c}^\top \mathbf{X}_S(\mathbf{X}_S^\top \mathbf{X}_S)^{-1} \dot{\rho}(\mathbf{v}_S + \hat{\beta}_S^o; \lambda)\|_\infty / \lambda^* : \|\mathbf{v}_S\|_\infty \leq \theta_1\lambda^* \}, \end{aligned}$$

where $S = \text{supp}(\beta)$ and $\hat{\beta}^o$ is the oracle LSE in Definition 1 (d). Note that when $\dot{\rho}(\hat{\beta}_S^o; \lambda) = 0$, $\theta_1 = 0$ is attained with $\mathbf{v}_S = 0$ and consequently $\theta_2 = 0$.

Theorem 6. (i) Let $S = \text{supp}(\beta)$ and \mathbf{P}_S be the projection to the column space of \mathbf{X}_S . Suppose $\rho(t; \lambda)$ is left- and right-differentiable in $t > 0$ and $\|\mathbf{X}_{S^c}^\top \mathbf{P}_S^\perp \epsilon\|_\infty \leq \dot{\rho}(0+; \lambda)$. Then, the oracle LSE $\hat{\beta}^o$ satisfies (25) with $\nu = \|\dot{\rho}(\hat{\beta}_S^o; \lambda)\|^2$. If in addition (17) holds and $\nu = 0 = \theta(\hat{\beta}^o, \kappa)$ with a certain $\kappa < \kappa_-(m + |S|)$ and m in (22) or (23), then $\hat{\beta}^o$ is the global solution of (2).
(ii) Suppose $\dot{\rho}(t; \lambda)$ is uniformly continuous in t in the region $\cup_{j \in S} [\hat{\beta}_j^o - \theta_1, \hat{\beta}_j^o + \theta_1]$. Suppose

$$\text{sgn}(\hat{\beta}^o) = \text{sgn}(\beta), \quad \min_{j \in S} |\hat{\beta}_j^o| > \theta_1\lambda^*, \quad \lambda^* \geq \|\mathbf{X}^\top \mathbf{P}_S^\perp \epsilon/n\|_\infty / (1 - \theta_2)_+. \quad (29)$$

Then, there exists a local solution $\tilde{\beta}^o$ of (2) satisfying $\text{sgn}(\tilde{\beta}^o) = \text{sgn}(\beta)$ and $\|\tilde{\beta}^o - \hat{\beta}^o\|_\infty \leq \theta_1\lambda^*$. If in addition (17) holds and $\theta(\tilde{\beta}^o, \kappa) = 0$ with a certain $\kappa < \kappa_-(m + |S|)$ and m in (22) or (23). Then, $\tilde{\beta}^o$ is the global solution of (2).

Remark 14. (i) For the capped- ℓ_1 penalty $\rho(t; \lambda) = \min(\gamma\lambda^2/2, \lambda|t|)$, $\nu = 0 = \theta(\hat{\beta}^o, \kappa)$ for $\kappa \geq 2/\gamma$ when $\min_{j \in S} |\hat{\beta}_j^o| > \gamma\lambda$. For the MCP $\rho(t; \lambda) = \lambda \int_0^{|t|} (1 - x/(\gamma\lambda))_+ dx$, $\theta(\cdot, \kappa) = 0$ for $\kappa \geq 1/\gamma$. For the SCAD penalty $\rho(t; \lambda) = \lambda \int_0^{|t|} \min\{1, (1 - (x/\lambda - 1)/(\gamma - 1))_+\} dx$, $\theta(\cdot, \kappa) = 0$ for $\kappa \geq 1/(\gamma - 1)$. (ii) For the ℓ_1 penalty, $\dot{\rho}(\mathbf{b}) = \text{sgn}(\mathbf{b})$ so that (29) is identical to (9) for the Lasso selection consistency. For concave penalties, $|\dot{\rho}(t; \lambda)|$ is small for large $|t|$, so that $\{\theta_1, \theta_2\}$ are typically smaller than $\{\theta_1^*, \theta_2^*\}$ for strong signals. In such cases, (29) is much weaker than (9).

For a nonconvex penalties such that $\dot{\rho}(t; \lambda) = 0$ when $|t| > a_0\lambda$ for some constant $a_0 > 0$, we automatically have $\dot{\rho}(\hat{\beta}_S^o; \lambda) = 0$ when $\min_{j \in S} |\hat{\beta}_j^o| > a_0\lambda$, which implies that $\theta_1 = \theta_2 = 0$. This special case gives the following easier to interpret corollary as a direct consequence of Theorems 5 and 6.

Corollary 3. Let $S = \text{supp}(\beta)$ and \mathbf{P}_S be the projection to the column space of \mathbf{X}_S . Suppose $\rho(t; \lambda)$ is left- and right-differentiable in $t > 0$ and $\|\mathbf{X}_{S^c}^\top \mathbf{P}_S^\perp \boldsymbol{\varepsilon}/n\|_\infty \leq \dot{\rho}(0+; \lambda)$. If (17) holds and $\dot{\rho}(\hat{\beta}_S^o; \lambda) = 0$, and $\theta(\hat{\beta}^o, \kappa) = 0$ with a certain $\kappa < \kappa_-(m + |S|)$ and m in (22) or (23), then $\hat{\beta}^o$ is the global solution of (2). Moreover, for any other exact local solution $\tilde{\beta}$ of (2) that is sparse with $|\text{supp}(\tilde{\beta}) \setminus S| \leq m$, we have $\tilde{\beta} = \hat{\beta}^o$.

Consider the simple examples of the capped- ℓ_1 penalty and MCP. For the capped- ℓ_1 penalty $\rho(t; \lambda) = \min(\gamma\lambda^2/2, \lambda|t|)$, we pick a sufficiently large γ such that $\gamma > 2/\kappa_- (|S| + m)$ for the m in (22) or (23). This will be possible with $m \asymp |S|$ when $\kappa_-(m)$ is uniformly bounded away from zero for small $m(\ln p)/n$ and $|S|(\ln p)/n$ is even smaller. For the MCP $\rho(t; \lambda) = \lambda \int_0^{|t|} (1 - x/(\lambda\gamma))_+ dx$, we pick $\gamma > 1/\kappa_- (|S| + m)$ for the m in (22) or (23). If $\min_{j \in S} |\hat{\beta}_j^o| \geq \gamma\lambda$, then the conditions of Corollary 3 are automatically satisfied for both penalties when $\|\mathbf{X}_{S^c}^\top \mathbf{P}_S^\perp \boldsymbol{\varepsilon}/n\|_\infty < \lambda$ (which can always be satisfied with a sufficiently large choice of λ). It follows that in this case, $\hat{\beta}^o$ is the global solution of (2), and there is no other local solution with no more than m nonzero-elements out of S . The essential condition here is the null consistency (17), which is an ℓ_2 condition. Note that in view of Corollary 2, the RIF condition is not essential for the equality of the global and oracle solutions in these examples, both with finite $\gamma^* = \gamma/2$. A similar result hold for the SCAD penalty, with somewhat different constant factors. The requirement of $\min_{j \in S} |\hat{\beta}_j^o| \geq \gamma\lambda$ is natural, and it directly follows (with probability $1 - \delta$) from the condition of $\min_{j \in S} |\beta_j| > \gamma\lambda + \sigma(1 + \sqrt{2 \ln(|S|/\delta)}) \lambda_{\min}^{-1/2}(\mathbf{X}_S^\top \mathbf{X}_S)$ under Assumption 1.

4.5 Approximate Global Solutions

We have mentioned in Remark 13 that gradient algorithm from the Lasso may yield the global solution of (2) for general $\rho(t; \lambda)$ under a sparse convexity condition or its generalization. Here we provide sufficient conditions for this to happen. This is done via a notion of approximate global solution. Given $\nu \geq 0$ and $\mathbf{b} \in \mathbb{R}^p$, we say that a vector $\tilde{\beta} \in \mathbb{R}^p$ is a $\{\nu, \mathbf{b}\}$ approximate global solution of (2) if

$$\left[\frac{1}{2n} \|\mathbf{X}\tilde{\beta} - \mathbf{y}\|_2^2 + \|\rho(\tilde{\beta}; \lambda)\|_1 \right] - \left[\frac{1}{2n} \|\mathbf{X}\mathbf{b} - \mathbf{y}\|_2^2 + \|\rho(\mathbf{b}; \lambda)\|_1 \right] \leq \nu. \quad (30)$$

To align different penalties at the same threshold level, we assume throughout this subsection that λ^* depends on $\rho(t; \lambda)$ only through λ in (12), e.g. $\lambda^* = \lambda$.

One method to find sparse local solution is to find a local solution that is also an approximate global solution. This can be achieved with the following simple procedure. First, we find the Lasso solution $\widehat{\beta}^{(\ell_1)}$ of (7). The following theorem shows that it is a $\{\nu, \beta\}$ approximate global solution of (2) with a relatively small ν under proper conditions. Now we can start with this solution $\widehat{\beta}^{(\ell_1)}$ and use gradient descent to find a local solution $\widetilde{\beta}$ of (2) that is also an approximate global solution. The following theorem then shows that under appropriate conditions, this local solution is sparse. Therefore results from Subsections 4.2 and 4.4 can be applied to relate it to the true global solution of (2).

Theorem 7. *Consider a penalty functions $\rho(t; \lambda)$ with $\lambda = \lambda^*$ in (12). Suppose the η null consistency condition (17) for $\rho(t; \lambda)$ with $0 < \eta < 1$.*

- (i) *Suppose $m = O(|S|)$ in (24) or under the SRC in Remark 11 for the Lasso $\widehat{\beta}^{(\ell_1)}$ in (7). Then, the Lasso $\widehat{\beta}^{(\ell_1)}$ is a $\{\nu, \beta\}$ approximate global solution for the penalty $\rho(t; \lambda)$ with $\nu \lesssim \lambda^2 |S|$.*
- (ii) *Assume that $\rho(t; \lambda)$ is continuous at $t = 0$. Let $\widetilde{\beta}$ be an local solution of (2) that is also a $\{\nu, \beta\}$ approximate global solution. Let $\xi' = 2/(1 - \eta)$. Consider $t_0 > 0$ and integer $m_0 > 0$ such that $\{2\kappa_+(m_0)b/m_0\}^{1/2} + \|\mathbf{X}^\top \varepsilon/n\|_\infty < \inf_{0 < s < t_0} \dot{\rho}(s; \lambda)$, where $b = \xi' \max\{\nu, \Delta(a'_1 \lambda_1^*, |S|; \lambda)\}$ with $a'_1 := (1 + \eta)/\text{RIF}_1(\xi', S)$ and $\lambda_1^* := \sup_{t \geq 0} |\dot{\rho}(t; \lambda)|$. Then,*

$$\#\{j \notin S : \widetilde{\beta}_j \neq 0\} < \widetilde{m} := m_0 + \lfloor b/\rho(t_0; \lambda) \rfloor.$$

Remark 15. *If $\rho(t; \lambda)$ is concave in t , then $\lambda_1^* = \dot{\rho}(0+; \lambda)$ and $\inf_{0 < s < t_0} \dot{\rho}(s; \lambda)$ can be replaced by $\dot{\rho}(t_0; \lambda)$ for choosing (t_0, m_0) . Theorem 7 applies to the ℓ_1 , capped- ℓ_1 , MCP and SCAD penalties with $\lambda = \lambda^* \asymp \lambda_1^*$, but not to the bridge penalty for which $\lambda_1^* = \infty$.*

Theorem 7 shows that the ℓ_1 solution $\widehat{\beta}^{(\ell_1)}$ is $\{\nu, \beta\}$ approximately global optimal with $\nu = O(|S|(\lambda^*)^2)$ in (30), and that a local solution $\widetilde{\beta}$ which is also approximate global optimal is a sparse local solution. Thus, with $b = O((\lambda^*)^2 |S|)$ and $\rho(t_0; \lambda) \asymp (\lambda^*)^2 \asymp (\lambda_1^*)^2$, the local solution $\widetilde{\beta}$ obtained with gradient descent from $\widehat{\beta}^{(\ell_1)}$ is sparse with $\#\{j \notin S : \widetilde{\beta}_j \neq 0\} = O(|S|)$. Here we assume that a line-search is performed in the gradient descent procedure so that the objective function always decreases (and thus each step leads to an $\{\nu, \beta\}$ approximate global optimal solution). Now Remark 12 can be applied to this sparse local solution, providing suitable conditions for this solution to be identical to the global optimal solution. If $\min_{j \in S} |\beta_j| > C\lambda_{\text{univ}}$ for a sufficiently large C , Corollary 3 (or Theorems 5 plus Theorem 6) can be applied to identify this local solution as the oracle LSE (or penalized LSE) and the global solution.

It is worth pointing out results of this paper concerning the global solution can be applied under the null consistency condition. For a general penalty function, this requires the condition (18) to hold. Although this is an ℓ_2 condition, it isn't needed for either ℓ_1 or ℓ_0 penalty as pointed out in Remark 9. In fact, this condition is also not needed if we consider local solution obtained with more specific numerical procedures such as [42, 47] that lead to specific sparse local solutions with oracle properties. Nevertheless, it is useful to observe that if the extra condition (18) holds, then such a local solution is also the unique global solution, and it can be obtained via other numerical procedures.

5 Technical Proofs

We first prove the following two lemmas, which will be useful in the analysis.

Lemma 1. If $\widehat{\beta}$ is the global solution of (2), then $\|\mathbf{X}^\top(\mathbf{y} - \mathbf{X}\widehat{\beta})/n\|_\infty \leq \lambda^*$. In particular, $\|\mathbf{X}^\top \varepsilon/n\|_\infty \leq \eta\lambda^*$ under the η null consistency condition (17).

Proof. The optimality of $\widehat{\beta}$ implies

$$\|\mathbf{y} - \mathbf{X}\widehat{\beta}\|_2^2/(2n) + \rho(\widehat{\beta}_j; \lambda) \leq \|\mathbf{y} - \mathbf{X}\widehat{\beta} - \mathbf{x}_j t\|_2^2/(2n) + \rho(\widehat{\beta}_j + t; \lambda)$$

for all real t . Since $\rho(t; \lambda)$ is subadditive in t ,

$$t\mathbf{x}_j^\top(\mathbf{y} - \mathbf{X}\widehat{\beta})/n \leq t^2\|\mathbf{x}_j\|_2^2/(2n) + \rho(\widehat{\beta}_j + t; \lambda) - \rho(\widehat{\beta}_j; \lambda) \leq t^2/2 + \rho(t; \lambda).$$

Since t is arbitrary, we obtain the desired bound via the definition of λ^* in (12). \square

Lemma 2. Assume the null consistency condition (17) with $\eta \in (0, 1)$. Suppose $\widehat{\beta} \in \mathbb{R}^p$ satisfy

$$\|\mathbf{y} - \mathbf{X}\widehat{\beta}\|_2^2/(2n) + \|\rho(\widehat{\beta}; \lambda)\|_1 \leq \|\mathbf{y} - \mathbf{X}\beta\|_2^2/(2n) + \|\rho(\beta; \lambda)\|_1 + \nu$$

with a certain $\nu > 0$. Let $\Delta = \widehat{\beta} - \beta$, $\xi = (1 + \eta)/(1 - \eta)$, and $S = \text{supp}(\beta)$. Then,

$$\|\mathbf{X}\Delta\|_2^2/(2n) + \|\rho(\Delta_{S^c}; \lambda)\|_1 \leq \xi\|\rho(\Delta_S; \lambda)\|_1 + \nu/(1 - \eta).$$

Proof. From the condition of the lemma, we have

$$\begin{aligned} 0 &\leq \nu + \|\mathbf{y} - \mathbf{X}\beta\|_2^2/(2n) + \|\rho(\beta; \lambda)\|_1 - \|\mathbf{y} - \mathbf{X}\widehat{\beta}\|_2^2/(2n) - \|\rho(\widehat{\beta}; \lambda)\|_1 \\ &= \nu - \|\mathbf{X}\Delta\|_2^2/(2n) + \varepsilon^\top \mathbf{X}\Delta/n + \|\rho(\beta; \lambda)\|_1 - \|\rho(\beta + \Delta; \lambda)\|_1. \end{aligned}$$

By (17), $\|\varepsilon/\eta\|_2^2/(2n) \leq \|\varepsilon/\eta - t\mathbf{X}\Delta\|_2^2/(2n) + \|\rho(t\Delta; \lambda)\|_1$ for all $t > 0$, which can be written as

$$\varepsilon^\top \mathbf{X}\Delta/n \leq \eta t\|\mathbf{X}\Delta\|_2^2/(2n) + (\eta/t)\|\rho(t\Delta; \lambda)\|_1.$$

The above two displayed inequalities yield

$$(1 - \eta t)\|\mathbf{X}\Delta\|_2^2/(2n) - \nu \leq (\eta/t)\|\rho(t\Delta; \lambda)\|_1 + \|\rho(\beta; \lambda)\|_1 - \|\rho(\beta + \Delta; \lambda)\|_1. \quad (31)$$

Now let $t = 1$. It follows from (31), $\beta_{S^c} = 0$, and then the sub-additivity of $\rho(t; \lambda)$ that

$$\begin{aligned} (1 - \eta)\|\mathbf{X}\Delta\|_2^2/(2n) - \nu &\leq \eta\|\rho(\Delta; \lambda)\|_1 + \|\rho(\beta_S; \lambda)\|_1 - \|\rho(\beta_S + \Delta_S; \lambda)\|_1 - \|\rho(\Delta_{S^c}; \lambda)\|_1 \\ &\leq (\eta + 1)\|\rho(\Delta_S; \lambda)\|_1 + (\eta - 1)\|\rho(\Delta_{S^c}; \lambda)\|_1. \quad \square \end{aligned}$$

5.1 Proof of Proposition 1

Let $t > 0$. By (12), $\rho(t; \lambda) \geq t(\lambda^* - t/2) \geq t\lambda^*/2$ for $t \leq \lambda^*$. For $t > \lambda^*$, $\rho(t; \lambda) \geq \rho(\lambda^*; \lambda) \geq (\lambda^*)^2/2$. This gives the lower bound of $\rho(t; \lambda)$. Let t_0 be the minimizer in (12) in the sense of $x/2 + \rho(x; \lambda)/x \rightarrow \lambda^*$ as $x \rightarrow t_0$ (when t_0 is a discontinuity of $\rho(\cdot; \lambda)$) or $x = t_0$. Let $x > 0$ and $q = \lfloor t/x \rfloor$. Since $\rho(t; \lambda)$ is nondecreasing and subadditive in $t > 0$, we have

$$\rho(t; \lambda) \leq \rho(qx; \lambda) + \rho(t - qx; \lambda) \leq (q + 1)\rho(x; \lambda) \leq (t + x)\rho(x; \lambda)/x.$$

It follows that (let $x \rightarrow t_0$) $\rho(t; \lambda) \leq (t + t_0)(\lambda^* - t_0/2) \leq \max_{t' \geq 0}(t + t')(\lambda^* - t'/2) = \rho^*(t; \lambda)$. The bound for $\Delta(a, k; \lambda)$ follows similarly from (let $x \rightarrow t_0$)

$$\|\rho(\mathbf{b}; \lambda)\|_1 \leq \sum_{j: b_j \neq 0} (|b_j| + x)\rho(x; \lambda)/x \leq k(a + x)\rho(x; \lambda)/x \leq k(a + t_0)(\lambda^* - t_0/2) \leq k\rho^*(a; \lambda).$$

The fact that $\rho^*(a; \lambda) \leq \max(a; 2\lambda^*)\lambda^*$ can be verified by simple algebra. \square

5.2 Proof of Proposition 2

Let $f(t) = t/\rho(t; \lambda)$ and A be the index set of the $|S|$ largest $|u_j|$. Since $\rho(t; \lambda)$ is nondecreasing in $|t|$, $\|\rho(\mathbf{u}_{S^c}; \lambda)\|_1 < \xi \|\rho(\mathbf{u}_S; \lambda)\|_1$ implies $\|\rho(\mathbf{u}_{A^c}; \lambda)\|_1 < \xi \|\rho(\mathbf{u}_A; \lambda)\|_1$. Since $f(t)$ is nondecreasing in t ,

$$\|\mathbf{u}_{A^c}\|_1 \leq \|\rho(\mathbf{u}_{A^c}; \lambda)\|_1 f(\|\mathbf{u}_{A^c}\|_\infty) \leq \xi \|\rho(\mathbf{u}_A; \lambda)\|_1 f(\|\mathbf{u}_{A^c}\|_\infty) \leq \xi \|\mathbf{u}_A\|_1.$$

This implies (15). In the above derivation, the first inequality follows from the definition of $f(t)$ and $\|\mathbf{u}_{A^c}\|_\infty \geq |u_j|$ for all $j \in A^c$; the second inequality is due to the condition $\|\rho(\mathbf{u}_{A^c}; \lambda)\|_1 < \xi \|\rho(\mathbf{u}_A; \lambda)\|_1$; the third inequality follows from the definition of $f(t)$ and the condition $\|\mathbf{u}_{A^c}\|_\infty \leq |u_j|$ for all $j \in A$. \square

5.3 Proof of Proposition 3

Since the left-hand side of (17) is increasing in $\rho(t; \lambda)$, we assume without loss of generality that

$$\rho(t; \lambda) = \min(\lambda^2/2, \lambda|t|), \quad \lambda = (1 + \zeta_0)(\sigma/\eta)\lambda_0, \quad \lambda_0 = (1 + \sqrt{2 \ln(2p/\delta)})/\sqrt{n}.$$

Since $\|\mathbf{X}^\top \boldsymbol{\varepsilon}/n\|_\infty \leq \max_{|A|=1} \|\mathbf{P}_A \boldsymbol{\varepsilon}\|_2/\sqrt{n}$ and $\|\mathbf{P}_A \boldsymbol{\varepsilon}\|_2 \leq \|\boldsymbol{\varepsilon}\|_2$, Assumption 1 implies that

$$\|\mathbf{X}^\top \boldsymbol{\varepsilon}/n\|_\infty \leq \sigma \lambda_0, \quad \|\mathbf{P}_A \boldsymbol{\varepsilon}\|_2 \leq \sigma \min\left[\sqrt{|A|n\lambda_0}, \sqrt{2n}\right], \quad \forall A \subseteq \{1, \dots, p\}, \quad (32)$$

with at least probability $1 - \exp(-n(1 - 1/\sqrt{2})^2) - \sum_{k=1}^n \binom{p}{k} (\delta/(2p))^k \geq 2 - \delta_n - e^{\delta/2}$.

Let $A = \{j : |b_j| > \lambda/2\}$ and $k = |A|$. It suffices to consider the case where A and \mathbf{b} satisfy

$$\mathbf{X}_A \mathbf{b}_A = \mathbf{P}_A(\boldsymbol{\varepsilon}/\eta - \mathbf{X}_{A^c} \mathbf{b}_{A^c}), \quad \text{rank}(\mathbf{P}_A) = |A| = k \leq \frac{\|\boldsymbol{\varepsilon}/\eta\|_2^2/(2n)}{\lambda^2/2} \leq \frac{2}{(1 + \zeta_0)^2 \lambda_0^2},$$

since these conditions hold for the global minimum for (2) with $\mathbf{y} = \boldsymbol{\varepsilon}/\eta$ and the capped- ℓ_1 penalty.

Under these conditions, we have $\mathbf{X} \mathbf{b} = \mathbf{P}_A \boldsymbol{\varepsilon}/\eta + \mathbf{P}_A^\perp \mathbf{X}_{A^c} \mathbf{b}_{A^c}$ and

$$\begin{aligned} & \|\boldsymbol{\varepsilon}/\eta\|_2^2/(2n) - \|\boldsymbol{\varepsilon}/\eta - \mathbf{X} \mathbf{b}\|_2^2/(2n) - \|\rho(\mathbf{b}; \lambda)\|_1 \\ &= (\mathbf{X} \mathbf{b})^\top (\boldsymbol{\varepsilon}/\eta)/n - \|\mathbf{X} \mathbf{b}\|_2^2/(2n) - \|\rho(\mathbf{b}; \lambda)\|_1 \\ &= \|\mathbf{P}_A(\boldsymbol{\varepsilon}/\eta)\|_2^2/(2n) + (\mathbf{P}_A^\perp \mathbf{X}_{A^c} \mathbf{b}_{A^c})^\top (\boldsymbol{\varepsilon}/\eta)/n - \|\mathbf{P}_A^\perp \mathbf{X}_{A^c} \mathbf{b}_{A^c}\|_2^2/(2n) - \|\rho(\mathbf{b}; \lambda)\|_1 \\ &\leq \|\mathbf{P}_A(\boldsymbol{\varepsilon}/\eta)\|_2^2/(2n) + (\mathbf{X}_{A^c} \mathbf{b}_{A^c})^\top (\boldsymbol{\varepsilon}/\eta)/n - (\mathbf{X}_{A^c} \mathbf{b}_{A^c})^\top \mathbf{P}_A(\boldsymbol{\varepsilon}/\eta)/n - \|\rho(\mathbf{b}; \lambda)\|_1 \\ &\leq \lambda_0^2 k (\sigma/\eta)^2/2 + \lambda_0 (\sigma/\eta) \|\mathbf{b}_{A^c}\|_1 - (\mathbf{X}_{A^c} \mathbf{b}_{A^c})^\top \mathbf{P}_A(\boldsymbol{\varepsilon}/\eta)/n - \|\rho(\mathbf{b}; \lambda)\|_1 \\ &< \left(\frac{1}{1 + \zeta_0} - 1\right) \|\rho(\mathbf{b}; \lambda)\|_1 - (\mathbf{X}_{A^c} \mathbf{b}_{A^c})^\top \mathbf{P}_A(\boldsymbol{\varepsilon}/\eta)/n. \end{aligned} \quad (33)$$

In the above derivation, the second inequality uses (32) and the third uses the fact that $\|\rho(\mathbf{b}; \lambda)\|_1 = \lambda^2 k/2 + \lambda \|\mathbf{b}_{A^c}\|_1$ by the definition of A and $\lambda = (1 + \zeta_0)(\sigma/\eta)\lambda_0$.

It follows from the shifting inequality in [7, 41] that

$$\begin{aligned} \left| \boldsymbol{\varepsilon}^\top \mathbf{P}_A \mathbf{X}_{A^c} \mathbf{b}_{A^c} \right| &\leq \max_{B \cap A = \emptyset, |B|=k} \|\mathbf{X}_B^\top \mathbf{P}_A \boldsymbol{\varepsilon}\|_2 \left(\|\mathbf{b}_{A^c}\|_\infty k^{1/2} + \|\mathbf{b}_{A^c}\|_1/k^{1/2} \right) \\ &\leq \max_{B \cap A = \emptyset, |B|=k} \|\mathbf{X}_B^\top \mathbf{P}_A \boldsymbol{\varepsilon}\|_2 (\lambda k/2 + \|\mathbf{b}_{A^c}\|_1)/\sqrt{k}. \end{aligned}$$

In the above derivation, the first inequality uses the shifting inequality and the second uses the fact that $\|\mathbf{b}_{A^c}\|_\infty \leq \lambda/2$ due to the definition of A . It follows from (18) and $\|\mathbf{P}_A \boldsymbol{\varepsilon}\|_2 \leq \sigma \lambda_0 \sqrt{nk}$ of (32) that for all $|A| = |B| = k \leq 2/\{(1 + \zeta_0)^2 \lambda_0^2\}$ with $B \cap A = \emptyset$,

$$\|\mathbf{X}_B^\top \mathbf{P}_A \boldsymbol{\varepsilon}\|_2 \leq \lambda_{\max}^{1/2}(\mathbf{X}_B^\top \mathbf{P}_A \mathbf{X}_B) \|\mathbf{P}_A \boldsymbol{\varepsilon}\|_2 \leq (\sigma \lambda_0 \sqrt{nk})(\zeta_0 \sqrt{n}) = \sigma \zeta_0 \lambda_0 n \sqrt{k}.$$

Thus, by combining the above two displayed inequalities, we find

$$\left| \boldsymbol{\varepsilon}^\top \mathbf{P}_A \mathbf{X}_{A^c} \mathbf{b}_{A^c} / (\eta n) \right| \leq \frac{\sigma \zeta_0 \lambda_0 n \sqrt{k}}{\eta n} \left(\frac{\lambda k/2 + \|\mathbf{b}_{A^c}\|_1}{\sqrt{k}} \right) = (1 + \zeta_0)^{-1} \zeta_0 \|\rho(\mathbf{b}; \lambda)\|_1.$$

due to $\lambda = (1 + \zeta_0)(\sigma/\eta) \lambda_0$ and $\|\rho(\mathbf{b}; \lambda)\|_1 = \lambda^2 k/2 + \lambda \|\mathbf{b}_{A^c}\|_1$. This and (33) yield the null consistency condition (17).

It remains to prove that (18) is an ℓ_2 -regularity condition on \mathbf{X} . Suppose that the rows of \mathbf{X} are iid from $N(0, \boldsymbol{\Sigma})$. Let $\mathbf{N}_{k,m}$ denote a $k \times m$ matrix with iid $N(0, 1)$ entries. We may write $\mathbf{X}_B = \mathbf{N}_{n,p}(\boldsymbol{\Sigma}^{1/2})_{p \times B}$. Let $\mathbf{U}\mathbf{U}^\top$ and $\mathbf{V}\mathbf{D}\mathbf{W}^\top$ be the SVDs of \mathbf{P}_A and $(\boldsymbol{\Sigma}^{1/2})_{p \times B}$ respectively. For fixed $\{A, B\}$, the entries of the $k \times k$ matrix $\mathbf{U}^\top \mathbf{N}_{n,p} \mathbf{V}$ are uncorrelated $N(0, 1)$ variables, so that we can write $\mathbf{P}_A \mathbf{X}_B = \mathbf{U} \mathbf{N}_{k,k} \mathbf{V}^\top (\boldsymbol{\Sigma}^{1/2})_{p \times B}$. Thus, by Theorem II.13 of [12],

$$\begin{aligned} P\{\lambda_{\max}^{1/2}(\mathbf{X}_B^\top \mathbf{P}_A \mathbf{X}_B) > (2k^{1/2} + t)\lambda_{\max}^{1/2}(\boldsymbol{\Sigma})\} &\leq P\{\lambda_{\max}^{1/2}(\mathbf{N}_{k,k}^\top \mathbf{N}_{k,k}) > 2k^{1/2} + t\} \\ &\leq \Phi(-t) \leq e^{-t^2/2}/2, \quad t > 0, \end{aligned}$$

where $\Phi(t)$ is the $N(0, 1)$ distribution function. Since there are no more than $\binom{p}{k, k, p-2k}$ choices of \mathbf{P}_A with rank k and $|B| = k$ with $A \cap B = \emptyset$,

$$\max_{B \cap A = \emptyset, |B|=|A|=k} \lambda_{\max}^{1/2}(\mathbf{X}_B^\top \mathbf{P}_A \mathbf{X}_B) \leq \lambda_{\max}^{1/2}(\boldsymbol{\Sigma}) \left(2k^{1/2} + \sqrt{8k \ln(2p/\delta)} \right), \quad \forall 1 \leq k \leq n, \quad (34)$$

with probability no smaller than

$$1 - \frac{1}{2} \sum_{k=1}^n \left(\frac{\delta}{2p} \right)^{4k} \binom{p}{k, k, p-2k} \geq 1 - \frac{1}{2} \sum_{k=1}^n \frac{(\delta^2/(4p))^{2k}}{(k!)^2} \geq 1 - \delta^4/(16p^2).$$

In the event (34), we have that for all $|A| = |B| = k$ and $k(1 + \zeta_0)^2(1 + \sqrt{2 \ln(2p/\delta)})^2 \leq 2n$,

$$\lambda_{\max}^{1/2}(\mathbf{X}_B^\top \mathbf{P}_A \mathbf{X}_B/n) \leq \frac{\lambda_{\max}^{1/2}(\boldsymbol{\Sigma}) (2k^{1/2} + \sqrt{8k \ln(2p/\delta)})}{\{k^{1/2}(1 + \zeta_0)(1 + \sqrt{2 \ln(2p/\delta)})\}/\sqrt{2}} = \frac{\sqrt{8} \lambda_{\max}^{1/2}(\boldsymbol{\Sigma})}{1 + \zeta_0}.$$

This proves the desired result. \square

5.4 Proof of Theorem 1

Let $\boldsymbol{\Delta} = \widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}$. Lemma 2 (with $\nu = 0$) implies that

$$\|\mathbf{X} \boldsymbol{\Delta}\|_2^2/(2n) + \|\rho(\boldsymbol{\Delta}_{S^c}; \lambda)\|_1 \leq \xi \|\rho(\boldsymbol{\Delta}_S; \lambda)\|_1. \quad (35)$$

Thus, (14) gives

$$\|\boldsymbol{\Delta}\|_q \leq \|\mathbf{X}^\top \mathbf{X} \boldsymbol{\Delta}\|_\infty |S|^{1/q} / \{n \text{RIF}_q(\xi, S)\}. \quad (36)$$

It follows from Lemma 1 that $\|\mathbf{X}^\top (\mathbf{y} - \mathbf{X} \widehat{\boldsymbol{\beta}})/n\|_\infty \leq \lambda^*$ and $\|\mathbf{X}^\top \boldsymbol{\varepsilon}/n\|_\infty \leq \eta \lambda^*$. Thus, we have $\|\mathbf{X}^\top \mathbf{X} \boldsymbol{\Delta}/n\|_\infty = \|\mathbf{X}^\top (\mathbf{y} - \mathbf{X} \widehat{\boldsymbol{\beta}} - \boldsymbol{\varepsilon})/n\|_\infty \leq (1 + \eta) \lambda^*$. This and (36) yield (19).

Now by combining the definition of $\Delta(a, |S|; \lambda)$ and $\|\boldsymbol{\Delta}\|_1 \leq (1 + \eta) \lambda^* |S| / \text{RIF}_1(\xi, S)$, which follows from (19), we obtain an estimate of $\|\rho(\boldsymbol{\Delta}_S; \lambda)\|_1$ in (35), which leads to the first inequality in (20). The second inequality in (20) then follows from Proposition 1 and Remark 4. \square

5.5 Proof of Theorem 2

Let $\widehat{S}_1 = \{j \in \widehat{S} \setminus S : |\widehat{\beta}_j| \geq t_0\}$ and $\widehat{S}_2 = \{j \in \widehat{S} \setminus S : |\widehat{\beta}_j| < t_0\}$. As in the proof of (20), it follows from the ℓ_1 error bound (19) and the definition of $\Delta(a_1, |S|; \lambda)$ in (13) that $\|\rho(\mathbf{\Delta}_S; \lambda)\|_1 \leq \Delta(a_1, |S|; \lambda)$ with the given a_1 . Thus,

$$|\widehat{S}_1| \leq \|\rho(\mathbf{\Delta}_{S^c}; \lambda)\|_1 / \rho(t_0; \lambda) \leq \xi \|\rho(\mathbf{\Delta}_S; \lambda)\|_1 / \rho(t_0; \lambda) \leq \xi \Delta(a_1, |S|; \lambda) / \rho(t_0; \lambda). \quad (37)$$

Let $\lambda_2 > \sqrt{2\xi\kappa_+(m_0)\Delta(a_1\lambda^*, |S|; \lambda)/m_0}$ satisfying $\lambda_2 + \|\mathbf{X}^\top \boldsymbol{\varepsilon}/n\|_\infty \leq \inf_{0 < s < t_0} \dot{\rho}(s; \lambda)$. The first order optimality condition implies that for all $j \in \widehat{S}$, $\mathbf{x}_j^\top (\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\beta}})/n = \dot{\rho}(t; \lambda)|_{t=\widehat{\beta}_j}$. For $j \in \widehat{S}_2$, $|\widehat{\beta}_j| \in (0, t_0)$, so that $|\mathbf{x}_j^\top (\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\beta}})/n| \geq (\lambda_2 + \|\mathbf{X}^\top \boldsymbol{\varepsilon}/n\|_\infty)$ by (21). Thus, for any set $A \subset \widehat{S}_2$ with $|A| \leq m_0$,

$$(\lambda_2 + \|\mathbf{X}^\top \boldsymbol{\varepsilon}/n\|_\infty)|A| \leq \|\mathbf{X}_A^\top (\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\beta}})/n\|_1 \leq \|\mathbf{X}_A^\top \boldsymbol{\varepsilon}/n\|_\infty |A| + |A|^{1/2} \|\mathbf{X}_A/\sqrt{n}\|_2 \|\mathbf{X}\mathbf{\Delta}\|_2/\sqrt{n}.$$

Since $\|\mathbf{X}_A/\sqrt{n}\|_2^2 \leq \kappa_+(m_0)$, $\lambda_2|A| \leq |A|^{1/2} \sqrt{\kappa_+(m_0)} \|\mathbf{X}\mathbf{\Delta}\|_2/\sqrt{n}$. It follows from Theorem 1 that $|A| \leq \kappa_+(m_0) \|\mathbf{X}\mathbf{\Delta}\|_2^2 / (n\lambda_2^2) \leq 2\xi\kappa_+(m_0)\Delta(a_1, |S|; \lambda)/\lambda_2^2 < m_0$. Thus, $\max_{A \subset \widehat{S}_2, |A| \leq m_0} |A| < m_0$, which implies that $|\widehat{S}_2| < m_0$. Combine this estimate with (37), we obtain the desired bound. \square

5.6 Proof of Theorem 3

It follows from the assumption of the theorem that for all $\mathbf{b} \in \mathbb{R}^p$,

$$\|\mathbf{X}\mathbf{b} - \boldsymbol{\varepsilon}/\eta\|_2^2 + \lambda^2 n \|\mathbf{b}\|_0 - \|\boldsymbol{\varepsilon}/\eta\|_2^2 = \|\mathbf{X}\mathbf{b}\|_2^2 + (2/\eta) \boldsymbol{\varepsilon}^\top \mathbf{X}\mathbf{b} + \lambda^2 n \|\mathbf{b}\|_0$$

is bounded from below by $\|\mathbf{X}\mathbf{b}\|_2^2 - 2\lambda\sqrt{n\|\mathbf{b}\|_0} \|\mathbf{X}\mathbf{b}\|_2 + \lambda^2 n \|\mathbf{b}\|_0 = (\|\mathbf{X}\mathbf{b}\|_2^2 - \lambda\sqrt{n\|\mathbf{b}\|_0})^2 \geq 0$. This implies the null-consistency condition. Moreover, (31) with $t = 1/\eta$ and $\nu = 0$ implies that

$$\|\widehat{\boldsymbol{\beta}}^{(\ell_0)}\|_0 - \|\boldsymbol{\beta}\|_0 \leq \eta^2 \|\widehat{\boldsymbol{\beta}}^{(\ell_0)} - \boldsymbol{\beta}\|_0 \leq \eta^2 \|\widehat{\boldsymbol{\beta}}^{(\ell_0)}\|_0 + \eta^2 \|\boldsymbol{\beta}\|_0,$$

which leads to the first bound of the theorem. The second bound is a direct consequence of Theorem 1, since $\Delta(\xi, |S|; \lambda) = \lambda^2 |S|/2$ by (13). \square

5.7 Proof of Theorem 4

For simplicity, let $\widehat{\boldsymbol{\beta}} = \widehat{\boldsymbol{\beta}}^{(\ell_0)}$, $\widehat{S} = \text{supp}(\widehat{\boldsymbol{\beta}})$, and $S = \text{supp}(\boldsymbol{\beta})$. We know that $\|\widehat{\boldsymbol{\beta}}\|_0 \leq (1 + \eta^2)/(1 - \eta^2) \|\boldsymbol{\beta}\|_0$ and thus $\|\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^o\|_0 \leq s$. Similar to the proof of Theorem 3, we have

$$\begin{aligned} 0 &\geq \|\mathbf{X}(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^o)\|_2^2 + 2(\mathbf{X}\widehat{\boldsymbol{\beta}}^o - \mathbf{y})^\top \mathbf{X}(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^o) + \lambda^2 n [\|\widehat{\boldsymbol{\beta}}\|_0 - \|\widehat{\boldsymbol{\beta}}^o\|_0] \\ &\geq \kappa_-(s)n \|\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^o\|_2^2 - \sqrt{2\kappa_-(s)\lambda n} \|(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^o)_{\widehat{S}-S}\|_1 + \lambda^2 n [\|\widehat{\boldsymbol{\beta}}\|_0 - \|\widehat{\boldsymbol{\beta}}^o\|_0] \\ &\geq \kappa_-(s)n \|(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^o)_S\|_2^2 + \kappa_-(s)n \|(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^o)_{\widehat{S}-S}\|_2^2 - 2\sqrt{0.5\lambda^2 n |\widehat{S} - S| \sqrt{\kappa_-(s)n}} \|(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}^o)_{\widehat{S}-S}\|_2 \\ &\quad + \lambda^2 n [\|\widehat{\boldsymbol{\beta}}\|_0 - \|\widehat{\boldsymbol{\beta}}^o\|_0] \\ &\geq \kappa_-(s)n \|(\widehat{\boldsymbol{\beta}}^o)_{S-\widehat{S}}\|_2^2 - 0.5\lambda^2 n |\widehat{S} - S| + \lambda^2 n [\|\widehat{\boldsymbol{\beta}}\|_0 - \|\widehat{\boldsymbol{\beta}}^o\|_0] \\ &\geq 2\lambda^2 n (|S - \widehat{S}| - \delta^o) - 0.5\lambda^2 n |\widehat{S} - S| + \lambda^2 n [\|\widehat{\boldsymbol{\beta}}\|_0 - \|\widehat{\boldsymbol{\beta}}^o\|_0] \\ &\geq \lambda^2 n (|S - \widehat{S}| + 0.5|\widehat{S} - S| - 2\delta^o). \end{aligned}$$

The first inequality uses the same derivation of a similar result in the proof of Theorem 3. The second inequality uses the assumption of the theorem, $(\mathbf{P}_S \boldsymbol{\varepsilon} - \boldsymbol{\varepsilon})^\top \mathbf{X} = (\mathbf{X} \hat{\boldsymbol{\beta}}^o - \mathbf{y})^\top \mathbf{X}$, and the fact that $(\mathbf{X} \hat{\boldsymbol{\beta}}^o - \mathbf{y})^\top \mathbf{X}_S = 0$. The forth inequality uses $b^2 - 2ab \geq -a^2$ and $\|(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)_S\|_2 \geq \|(\hat{\boldsymbol{\beta}}^o)_{S-\hat{S}}\|_2$. The fifth inequality uses

$$\begin{aligned} \kappa_-(s)n\|(\hat{\boldsymbol{\beta}}^o)_{S-\hat{S}}\|_2^2 &\geq \kappa_-(s)n \sum_{j \in S-\hat{S}; |\hat{\beta}_j|^2 \geq 2\lambda/\kappa_-(s)} (\hat{\beta}_j^o)^2 \\ &\geq 2\lambda^2 n \left| \{j \in S-\hat{S}; (\hat{\beta}_j^o)^2 \geq 2\lambda^2/\kappa_-(s)\} \right| \geq 2\lambda^2 n(|S-\hat{S}| - \delta^o). \end{aligned}$$

The last inequality uses the derivation $\|\hat{\boldsymbol{\beta}}\|_0 - \|\hat{\boldsymbol{\beta}}^o\|_0 \geq |\hat{S}| - |S| = |\hat{S} - S| - |S - \hat{S}|$ and simple algebra. This proves the first desired bound. Similarly, we have

$$\begin{aligned} 0 &\geq \|\mathbf{X}(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)\|_2^2 - \sqrt{2\kappa_-(s)\lambda}n\|(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)_{\hat{S}-S}\|_1 + \lambda^2 n[\|\hat{\boldsymbol{\beta}}\|_0 - \|\hat{\boldsymbol{\beta}}^o\|_0] \\ &\geq 0.5\|\mathbf{X}(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)\|_2^2 + 0.5\kappa_-(s)n\|(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)_S\|_2^2 + 0.5\kappa_-(s)n\|(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)_{\hat{S}-S}\|_2^2 \\ &\quad - \sqrt{2\kappa_-(s)|\hat{S}-S|\lambda}n\|(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)_{\hat{S}-S}\|_2 + \lambda^2 n[\|\hat{\boldsymbol{\beta}}\|_0 - \|\hat{\boldsymbol{\beta}}^o\|_0] \\ &\geq 0.5\|\mathbf{X}(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)\|_2^2 + 0.5\kappa_-(s)n\|(\hat{\boldsymbol{\beta}}^o)_{S-\hat{S}}\|_2^2 - \lambda^2 n|\hat{S}-S| + \lambda^2 n[\|\hat{\boldsymbol{\beta}}\|_0 - \|\hat{\boldsymbol{\beta}}^o\|_0] \\ &\geq 0.5\|\mathbf{X}(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)\|_2^2 + \lambda^2 n(|S-\hat{S}| - \delta^o) - \lambda^2 n|\hat{S}-S| + \lambda^2 n[\|\hat{\boldsymbol{\beta}}\|_0 - \|\hat{\boldsymbol{\beta}}^o\|_0] \\ &\geq 0.5\|\mathbf{X}(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)\|_2^2 - \lambda^2 n\delta^o. \end{aligned}$$

The second inequality uses the definition of $\kappa_-(s)$. The third inequality uses $0.5b^2 - \sqrt{2}ab \geq -a^2$ and $\|(\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^o)_S\|_2 \geq \|(\hat{\boldsymbol{\beta}}^o)_{S-\hat{S}}\|_2$. The fourth inequality uses the previously derived inequality $\kappa_-(s)n\|(\hat{\boldsymbol{\beta}}^o)_{S-\hat{S}}\|_2^2 \geq 2\lambda^2 n(|S-\hat{S}| - \delta^o)$. The last inequality uses the derivation $\|\hat{\boldsymbol{\beta}}\|_0 - \|\hat{\boldsymbol{\beta}}^o\|_0 \geq |\hat{S}| - |S| = |\hat{S} - S| - |S - \hat{S}|$ and simple algebra. This leads to the second desired bound. \square

5.8 Proof of Theorem 5

Since $\tilde{\boldsymbol{\beta}}^{(j)}$ are approximate local solutions with excess $\nu^{(j)}$, (25) gives

$$\|\mathbf{X}^\top \mathbf{X} \boldsymbol{\Delta} / n + \dot{\rho}(\tilde{\boldsymbol{\beta}}^{(1)}; \lambda) - \dot{\rho}(\tilde{\boldsymbol{\beta}}^{(2)}; \lambda)\|_2 \leq (\nu^{(1)})^{1/2} + (\nu^{(2)})^{1/2} \leq \sqrt{\nu}.$$

Let $E = \tilde{S}^{(1)} \cup \tilde{S}^{(2)}$. Since $|E| \leq m+k$ and $\kappa < \kappa_-(m+k)$, it follows that

$$\begin{aligned} \|\mathbf{X} \boldsymbol{\Delta}\|_2^2 / n &\leq -\boldsymbol{\Delta}^\top (\dot{\rho}(\tilde{\boldsymbol{\beta}}^{(1)}; \lambda) - \dot{\rho}(\tilde{\boldsymbol{\beta}}^{(2)}; \lambda)) + \sqrt{\nu} \|\boldsymbol{\Delta}\|_2 \\ &\leq \kappa \|\boldsymbol{\Delta}\|_2^2 + \left| (\boldsymbol{\Delta}^\top \theta(|\tilde{\boldsymbol{\beta}}^{(1)}|, \kappa)) \right| + \sqrt{\nu} \|\boldsymbol{\Delta}\|_2 \\ &\leq \kappa \|\boldsymbol{\Delta}\|_2^2 + \left(\|\theta(|\tilde{\boldsymbol{\beta}}_E^{(1)}|, \kappa)\|_2 + \sqrt{\nu} \right) \|\boldsymbol{\Delta}\|_2 \end{aligned}$$

Since $\|\boldsymbol{\Delta}\|_2^2 \leq \|\mathbf{X} \boldsymbol{\Delta}\|_2^2 / \{n\kappa_-(m+k)\}$, (26) follows.

Let $E_1 := \{j : |\tilde{\beta}_j^{(1)} - \tilde{\beta}_j^{(2)}| \geq \lambda_0 / \sqrt{\kappa_-(m+k)}\}$. We have $\lambda_0^2 |E_1| \leq \kappa_-(m+k) \|\boldsymbol{\Delta}\|_2^2 \leq \|\mathbf{X} \boldsymbol{\Delta}\|_2^2 / n$.

Since $j \in S \setminus \tilde{S}^{(2)}$ implies $\tilde{\beta}_j^{(1)} - \tilde{\beta}_j^{(2)} = \tilde{\beta}_j^{(1)}$, (27) follows.

Let $E_2 := \tilde{S}^{(2)} \setminus S$ and $\lambda'_0 = \dot{\rho}(0+; \lambda) - \|\mathbf{X}_{S^c}^\top (\mathbf{X} \tilde{\boldsymbol{\beta}}^{(1)} - \mathbf{y}) / n\|_\infty$. For $j \in E_2$,

$$\lambda'_0 \leq \dot{\rho}(0+; \lambda) + \text{sgn}(\tilde{\beta}^{(2)}) x_j^\top (\mathbf{X} \tilde{\boldsymbol{\beta}}^{(1)} - \mathbf{y}) / n$$

$$\leq \{ \dot{\rho}(0+; \lambda) - \text{sgn}(\tilde{\beta}^{(2)}) \dot{\rho}(\tilde{\beta}_j^{(2)}; \lambda) \} + |\mathbf{x}_j^\top (\mathbf{X} \tilde{\beta}^{(2)} - \mathbf{y})/n + \dot{\rho}(\tilde{\beta}_j^{(2)}; \lambda)| + |\mathbf{x}_j^\top \mathbf{X} \Delta/n|.$$

Since $\theta(0+, \kappa) = 0$ means $\dot{\rho}(0+; \lambda) - \text{sgn}(t) \dot{\rho}(t; \lambda) = \dot{\rho}(0+; \lambda) - \dot{\rho}(|t|; \lambda) \leq \kappa|t|$ for $t \neq 0$ and $\tilde{\beta}_{E_2}^{(2)} = -\Delta_{E_2}$,

$$\begin{aligned} |E_2| \lambda'_0 &\leq \kappa \|\Delta_{E_2}\|_1 + \|\mathbf{X}_{E_2}^\top (\mathbf{X} \tilde{\beta}^{(2)} - \mathbf{y})/n + \dot{\rho}(\tilde{\beta}_{E_2}^{(2)}; \lambda)\|_1 + \|\mathbf{X}_{E_2}^\top \mathbf{X} \Delta/n\|_1 \\ &\leq \sqrt{|E_2|} \left\{ \kappa \|\Delta\|_2 + \sqrt{\tilde{\nu}^{(2)}} + \|\mathbf{X}_{E_2}^\top \mathbf{X} \Delta/n\|_2 \right\} \end{aligned}$$

Since $\|\Delta\|_2^2 \leq \|\mathbf{X} \Delta\|_2^2 / \{n\kappa_-(m+k)\}$ and $\|\mathbf{X}_{E_2}^\top \mathbf{X} \Delta/n\|_2^2 \leq \kappa_+(m) \|\mathbf{X} \Delta\|_2^2 / n$, (28) follows. \square

5.9 Proof of Theorem 6

We note that $\mathbf{X}^\top (\mathbf{y} - \mathbf{X} \hat{\beta}^o) = \mathbf{X}^\top (\boldsymbol{\varepsilon} - \mathbf{P}_S \boldsymbol{\varepsilon}) = \mathbf{X}^\top \mathbf{P}_S^\perp \boldsymbol{\varepsilon}$.

(i) Since $\mathbf{x}_j^\top (\mathbf{X} \hat{\beta}^o - \mathbf{y})/n + \dot{\rho}(\hat{\beta}_j^o; \lambda) = \dot{\rho}(\hat{\beta}_j^o; \lambda)$ for $j \in S$ and $\mathbf{x}_j^\top (\mathbf{X} \hat{\beta}^o - \mathbf{y})/n + \dot{\rho}(\hat{\beta}_j^o; \lambda) = 0$ for $j \notin S$, $\nu = \|\dot{\rho}(\hat{\beta}_S^o; \lambda)\|^2$. Let $\hat{\beta}^{(2)}$ be the global solution of (2). Under the additional conditions, $\hat{\beta}^{(2)} = \hat{\beta}^o$ by Theorems 2 and (26) with $\hat{\beta}^{(1)} = \hat{\beta}^o$.

(ii) Since $\min_{j \in S} |\hat{\beta}_j^o| \geq \theta_1 \lambda^*$, the map $\mathbf{b}_S \rightarrow \hat{\beta}_S^o - (\mathbf{X}_S^\top \mathbf{X}_S/n)^{-1} \dot{\rho}(\mathbf{b}_S; \lambda)$ is continuous and closed in the rectangle $B = \{\mathbf{v} : \|\mathbf{v}_S - \hat{\beta}_S^o\|_\infty \leq \theta_1 \lambda^*, \mathbf{v}_{S^c} = 0\}$. Thus, the Brouwer fixed point theorem implies a $\hat{\beta} \in B$ satisfying $\text{sgn}(\hat{\beta}) = \text{sgn}(\hat{\beta}^o)$ and

$$\mathbf{X}_S^\top (\mathbf{y} - \mathbf{X} \hat{\beta}) = (\mathbf{X}_S^\top \mathbf{X}_S/n) (\hat{\beta}^o - \hat{\beta})_S = \dot{\rho}(\hat{\beta}_S; \lambda).$$

Since $\mathbf{y} - \mathbf{X} \hat{\beta} = \mathbf{y} - \mathbf{X} \hat{\beta}^o - \mathbf{X}_S (\hat{\beta} - \hat{\beta}^o)_S = \mathbf{y} - \mathbf{X} \hat{\beta}^o - \mathbf{X}_S (\mathbf{X}_S^\top \mathbf{X}_S/n)^{-1} \dot{\rho}(\hat{\beta}_S; \lambda)$,

$$\|\mathbf{X}_{S^c}^\top (\mathbf{y} - \mathbf{X} \hat{\beta})/n\|_\infty \leq \|\mathbf{X}_{S^c}^\top (\mathbf{y} - \mathbf{X} \hat{\beta}^o)/n\|_\infty + \theta_2 \lambda^* \leq \lambda^*,$$

so that $\hat{\beta}$ is a local solution of (2). The proof of global optimality of $\tilde{\beta}^o$ is the same as (i). \square

5.10 Proof of Theorem 7

The proof is similar to that of Theorem 2. As intermediate results, we will prove lemmas that are analogous to Lemma 1 and Theorem 1. In the following, we assume that the conditions of the theorem hold. We also let $\Delta = \tilde{\beta} - \beta$.

Lemma 3. Let $\lambda_1^* := \sup_{t \geq 0} |\dot{\rho}(t; \lambda)|$. We have $\|\mathbf{X}^\top (\mathbf{y} - \mathbf{X} \tilde{\beta})/n\|_\infty \leq \lambda_1^*$.

Proof. A local solution satisfies $|\mathbf{x}_j^\top (\mathbf{X} \tilde{\beta} - \mathbf{y})/n| = |\dot{\rho}(\tilde{\beta}_j; \lambda)| \leq \lambda_1^*$ for all j . \square

Lemma 4. We have $\|\mathbf{X} \Delta\|_2^2 / (2n) + \|\rho(\Delta_{S^c}; \lambda)\|_1 \leq b$ with $\Delta = \tilde{\beta} - \beta$.

Proof. We consider two situations: the first is $\|\rho(\Delta_S; \lambda)\|_1 \leq \nu$, and the second is $\|\rho(\Delta_S; \lambda)\|_1 > \nu$. In the first situation, we obtain directly from Lemma 2 that

$$\|\mathbf{X} \Delta\|_2^2 / (2n) + \|\rho(\Delta_{S^c}; \lambda)\|_1 \leq 2\nu / (1 - \eta) = \xi' \nu.$$

In the second situation, we obtain from Lemma 2 that $\|\mathbf{X} \Delta\|_2^2 / (2n) + \|\rho(\Delta_{S^c}; \lambda)\|_1 \leq \xi' \|\rho(\Delta_S; \lambda)\|_1$. Therefore (14) gives $\|\Delta\|_1 \leq \|\mathbf{X}^\top \mathbf{X} \Delta\|_\infty |S| / \{n \text{RIF}_1(\xi', S)\}$. It follows from Lemma 3 that

$\|\mathbf{X}^\top(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})/n\|_\infty \leq \lambda_1^*$. Similarly, $\|\mathbf{X}^\top(\boldsymbol{\varepsilon}/\eta)/n\|_\infty \leq \lambda^* = \lambda$ due to (17). Since $\lambda = \lambda^* \leq \inf_t |\rho(t; \lambda)/t| \leq \lambda_1^*$, we have $\|\mathbf{X}^\top \mathbf{X} \boldsymbol{\Delta}/n\|_\infty = \|\mathbf{X}^\top(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}} - \boldsymbol{\varepsilon})/n\|_\infty \leq (1 + \eta)\lambda_1^*$. This implies that $\|\boldsymbol{\Delta}\|_1 \leq a'_1 \lambda_1^* |S|$, where $a'_1 = (1 + \eta)/\text{RIF}_1(\xi', S)$. This can be combined with Lemma 2 and the definition of $\Delta(a, |S|; \lambda)$ to obtain

$$\|\mathbf{X} \boldsymbol{\Delta}\|_2^2/(2n) + \|\rho(\boldsymbol{\Delta}_{S^c}; \lambda)\|_1 \leq \xi' \Delta(a'_1 \lambda_1^*, |S|; \lambda).$$

Combine the two situations, we obtain the lemma. \square

We are now ready to prove the theorem.

(i) Let $\boldsymbol{\Delta}^{(\ell_1)} = \hat{\boldsymbol{\beta}}^{(\ell_1)} - \boldsymbol{\beta}$. Since $|\boldsymbol{\varepsilon}^\top \mathbf{X} \boldsymbol{\Delta}^{(\ell_1)}|/n \leq \|\mathbf{X} \boldsymbol{\Delta}^{(\ell_1)}\|_2^2/(2n) + \|\rho(\boldsymbol{\Delta}^{(\ell_1)}; \lambda)\|_1$ by (17),

$$\begin{aligned} \nu &= \|\mathbf{X} \boldsymbol{\Delta}^{(\ell_1)}\|_2^2/(2n) - \boldsymbol{\varepsilon}^\top \mathbf{X} \boldsymbol{\Delta}^{(\ell_1)}/n + \|\rho(\hat{\boldsymbol{\beta}}^{(\ell_1)}; \lambda)\|_1 - \|\rho(\boldsymbol{\beta}; \lambda)\|_1 \\ &\leq 2 \left\{ \|\mathbf{X} \boldsymbol{\Delta}^{(\ell_1)}\|_2^2/(2n) + \|\rho(\boldsymbol{\Delta}^{(\ell_1)}; \lambda)\|_1 \right\} \\ &\leq 2 \left\{ \xi a_1 \lambda^2 |S| + \Delta(a_1 \lambda |S|/m, m; \lambda) \right\} \\ &\leq 2 \left\{ \xi a_1 \lambda^2 |S| + \lambda^2 \max(a_1 |S|, 2m) \right\} = O(\lambda^2 |S|). \end{aligned}$$

(ii) Let $\hat{S}_1 = \{j \in \hat{S} \setminus S : |\hat{\beta}_j| \geq t_0\}$, $\hat{S}_2 = \{j \in \hat{S} \setminus S : |\hat{\beta}_j| < t_0\}$, and $\lambda_2 > \sqrt{2\kappa_+(m_0)b/m_0}$ satisfying $\lambda_2 + \|\mathbf{X}^\top \boldsymbol{\varepsilon}/n\|_\infty < \inf_{0 < s < t_0} \dot{\rho}(s; \lambda)$. Just as in the proof of Theorem 2, we have $|\hat{S}_1| \leq \|\rho(\boldsymbol{\Delta}_{S^c}; \lambda)\|_1/\rho(t_0; \lambda)$, and for any $A \subset \hat{S}_2$ with $|A| \leq m_0$, $|A| \leq \kappa_+(m_0)\|\mathbf{X} \boldsymbol{\Delta}\|_2^2/(n\lambda_2^2)$. We apply Lemma 4 to obtain $|\hat{S}_1| \leq b/\rho(t_0; \lambda)$ and $|A| \leq 2\kappa_+(m_0)b/\lambda_2^2 < m_0$. Thus, $\max_{A \subset \hat{S}_2, |A| \leq m_0} |A| < m_0$, which implies that $|\hat{S}_2| < m_0$. The theorem follows. \square

References

- [1] H. Akaike. Information theory and an extension of the maximum likelihood principle. In *Second International Symposium on Information Theory*, pages 267–281. Akademiai Kiadó, Budapest, 1973.
- [2] A. Antoniadis. Comments on: ℓ_1 -penalization for mixture regression models. *TEST*, 19:257–258, 2010.
- [3] P. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *Annals of Statistics*, 37:1705–1732, 2009.
- [4] P. Breheny and J. Huang. Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Annals of Applied Statistics*, 5:232–253, 2011.
- [5] P. Bühlmann and S. van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer, New York, 2011.
- [6] F. Bunea, A. Tsybakov, and M. Wegkamp. Sparsity oracle inequalities for the lasso. *Electronic Journal of Statistics*, 1:169–194, 2007.
- [7] T. Cai, L. Wang, and G. Xu. Shifting inequality and recovery of sparse signals. *IEEE Transactions on Signal Processing*, 58:1300–1308, 2010.

- [8] E. Candes and Y. Plan. Near-ideal model selection by ℓ_1 minimization. *Annals of Statistics*, 37:2145–2177, 2009.
- [9] E. Candes and T. Tao. The Dantzig selector: statistical estimation when p is much larger than n (with discussion). *Annals of Statistics*, 35:2313–2404, 2007.
- [10] E. J. Candes and T. Tao. Decoding by linear programming. *IEEE Trans. on Information Theory*, 51:4203–4215, 2005.
- [11] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Review*, 43:129–159, 2001.
- [12] K. Davidson and S. Szarek. Local operator theory, random matrices and Banach spaces. In *Handbook on the Geometry of Banach Spaces*, volume 1. 2001.
- [13] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression (with discussion). *Annals of Statistics*, 32:407–499, 2004.
- [14] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96:1348–1360, 2001.
- [15] J. Fan and H. Peng. On non-concave penalized likelihood with diverging number of parameters. *Annals of Statistics*, 32:928–961, 2004.
- [16] I. Frank and J. Friedman. A statistical view of some chemometrics regression tools (with discussion). *Technometrics*, 35:109–148, 1993.
- [17] E. Greenshtein and Y. Rotiv. Persistence in high-dimensional linear predictor selection and the virtue of overparametrization. *Bernoulli*, 10:971–988, 2004.
- [18] J. Huang, S. Ma, and C.-H. Zhang. Adaptive lasso for sparse high-dimensional regression models. *Statistica Sinica*, 18:1603–1618, 2008.
- [19] D. Hunter and R. Li. Variable selection using MM algorithms. *Annals of Statistics*, 33:1617–1642, 2005.
- [20] Y. Kim, H. Choi, and H.-S. Oh. Smoothly clipped absolute deviation on high dimensions. *Journal of American Statistical Association*, 103:1665–1673, 2008.
- [21] K. Knight and W. Fu. Asymptotics for lasso-type estimators. *Annals of Statistics*, 28:1356–1378, 2000.
- [22] V. Koltchinskii. The dantzig selector and sparsity oracle inequalities. *Bernoulli*, 15:799–828, 2009.
- [23] J. Liu, P. Wonka, and J. Ye. Multi-stage Dantzig selector. In *NIPS 10*. 2010.
- [24] C. Mallows. Some comments on cp. *Technometrics*, 12:661–675, 1973.
- [25] R. Mazumder, J. Friedman, and T. Hastie. Sparsenet : Coordinate descent with non-convex penalties. *Journal of American Statistical Association*, page in press, 2011.

- [26] N. Meinshausen and P. Bühlmann. High dimensional graphs and variable selection with the Lasso. *Annals of Statistics*, 34:1436–1462, 2006.
- [27] N. Meinshausen and B. Yu. Lasso-type recovery of sparse representations for high-dimensional data. *Annals of Statistics*, 37:246–270, 2009.
- [28] M. Osborne, B. Presnell, and B. Turlach. A new approach to variable selection in least squares problems. *IMA Journal of Numerical Analysis*, 20:389–404, 2000.
- [29] M. Osborne, B. Presnell, and B. Turlach. On the lasso and its dual. *Journal of Computational and Graphical Statistics*, 9(2):319–337, 2000.
- [30] G. Raskutti, M. J. Wainwright, and B. Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. Technical report, University of California, Berkeley, 2009.
- [31] G. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6:461–464, 1978.
- [32] N. Städler, P. Bühlmann, and S. van de Geer. ℓ_1 -penalization for mixture regression models (with discussion). *Test*, 2:209–285, 2010.
- [33] T. Sun and C.-H. Zhang. Comments on: ℓ_1 -penalization for mixture regression models. *Test*, 2:270–275, 2010.
- [34] T. Sun and C.-H. Zhang. Scaled sparse linear regression. Technical Report arXiv:1104.4595, arXiv, 2011.
- [35] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 58:267–288, 1996.
- [36] J. A. Tropp. Just relax: convex programming methods for identifying sparse signals in noise. *IEEE Transactions on Information Theory*, 52:1030–1051, 2006.
- [37] S. van de Geer. The deterministic Lasso. Technical Report 140, ETH Zurich, Switzerland, 2007.
- [38] S. van de Geer. High-dimensional generalized linear models and the Lasso. *Annals of Statistics*, 36:614–645, 2008.
- [39] S. van de Geer and P. Bühlmann. On the conditions used to prove oracle results for the Lasso. *Electronic Journal of Statistics*, 3:1360–1392, 2009.
- [40] M. J. Wainwright. Sharp thresholds for noisy and high-dimensional recovery of sparsity using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Transactions on Information Theory*, 55:2183–2202, 2009.
- [41] F. Ye and C.-H. Zhang. Rate minimaxity of the lasso and dantzig selector for the ℓ_q loss in ℓ_r balls. *Journal of Machine Learning Research*, 11:3519–3540, 2010.
- [42] C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38:894–942, 2010.

- [43] C.-H. Zhang and J. Huang. The sparsity and bias of the Lasso selection in high-dimensional linear regression. *Annals of Statistics*, 36:1567–1594, 2008.
- [44] T. Zhang. Some sharp performance bounds for least squares regression with l_1 regularization. *Annals of Statistics*, 37:2109–2144, 2009.
- [45] T. Zhang. Analysis of multi-stage convex relaxation for sparse regularization. *Journal of Machine Learning Research*, 11:1087–1107, 2010.
- [46] T. Zhang. Adaptive forward-backward greedy algorithm for learning sparse representations. *IEEE Transactions on Information Theory*, 57:4689–4708, 2011.
- [47] T. Zhang. Multi-stage convex relaxation for feature selection. Technical Report arXiv:1106.0565, arXiv, 2011.
- [48] P. Zhao and B. Yu. On model selection consistency of Lasso. *Journal of Machine Learning Research*, 7:2541–2567, 2006.
- [49] H. Zou. The adaptive Lasso and its oracle properties. *Journal of the American Statistical Association*, 101:1418–1429, 2006.
- [50] H. Zou and R. Li. One-step sparse estimates in nonconcave penalized likelihood models (with discussion). *Annals of Statistics*, 36(4):1509–1533, 2008.